Material Requirements Planning under uncertainty

Optimising safety stock levels of changeover materials using simulation-based optimisation

TIL Master Thesis

Gijs Janssen

Photo: © Kaspars Grinvalds / Adobe Stock [1]

Material Requirements Planning under uncertainty

Optimising safety stock levels of changeover materials using simulation-based optimisation

A Master Thesis

Author:
Gijs Janssen (5083478)

Research supervisors:
Dr. B. (Bilge) Atasoy (TU Delft)
Dr. J. (Jie) Gao (TU Delft)

Company supervisor:
Dr. B. (Ben) Hermans (ORTEC)

Master of Science

Transport, Infrastructure and Logistics

Faculty of Civil Engineering and Geosciences

Delft University of Technology

Report number: 2026.TIL.9221

Date: 15th September 2026

Acknowledgements

I would like to extend a thank you to my research supervisors, Bilge Atasoy and Jie Gao, for their guidance and support throughout this thesis project. Their critical feedback and continuous encouragement through frequent discussion sessions were really beneficial for me during the past months. I really enjoyed working with them again during the final stage of my master's degree.

I would also like to thank my company supervisor, Ben Hermans, for taking the time to supervise me at ORTEC during this thesis project. His knowledge on the application of this research was frequently consulted and it provided me with many directions and possibilities throughout this research. Besides his practical insights, his ongoing support and encouragement was appreciated greatly.

Finally, I would like to thank everyone who contributed indirectly to this research, whether at ORTEC or at TU Delft, whether through discussions, support or with anything else along the way.

Gijs Janssen
Delft, September 2026

Abstract

Material Requirements Planning (MRP) is widely used to translate production plans into procurement and inventory decisions. Its performance can deteriorate in environments characterised by demand uncertainty and frequent material changeovers. In such settings, deterministic MRP models often set safety stocks for changeover materials to zero in order to minimise material obsolescence. This increases vulnerability to stockouts, minimum-order-quantity rush orders and last-minute procurement costs. This thesis investigates whether simulation-based optimisation can be used to parametrise safety stock levels for changeover materials in a deterministic mixed-integer programming MRP model to improve rolling-horizon performance in a stochastic production environment.

A framework is developed in which the manufacturer's existing MRP optimisation model is treated as a black box and is embedded in an Optimisation with Simulation-based Iterations (OSI) structure. To better represent robustness in the MRP logic, a hard safety stock constraint is introduced through which safety stock levels are optimised for grouped materials. Demand uncertainty is modelled using historical weekly production plans, analysed through a method based on block bootstrapping and kernel density estimation (KDE). This analysis showed that nearly all observed variability in the production plans is demand-driven rather than induced by planning decisions.

The framework is tested using a Full Enumeration optimisation technique, and experimented on different simulation settings, material grouping strategies, uncertainty variants and safety stock variants. Although the proposed approach is technically feasible and yields a reproducible method for evaluating safety stock policies under uncertainty, the final optimisation does not deliver an obsolete stock costs improvement caused by the addition of safety stock. Low safety stock levels consistently performed best, suggesting that a generic safety stock policy for changeover materials is not effective under the current objective structure and framework setup. The main contribution of this thesis is methodological: it demonstrates how simulation-based optimisation can be applied to an industry-scale, black-box MRP environment. It also highlights the challenges of translating robustness into measurable system-wide gains when only obsolete stock costs are optimised. The findings indicate that future research should adopt broader objective functions and further refine material grouping and uncertainty scenarios to better capture material- and scenario-specific trade-offs.

Contents

List of Figures

List of Tables

Terminology

Abbreviation	Definition
ACO	Ant Colony Optimisation
ANN	Artificial Neural Network
API	Application Programming Interface
B2B	Business-to-Business
BOM	Bill of Materials
BO	Bayesian Optimisation
CMG	Common Material Group
DQN	Deep Q-Network
DOS	Days of Sales
EI	Expected Improvement
EOQ	Economic Order Quantity
ERP	Enterprise Resource Planning
FMCG	Fast-Moving Consumer Goods
FOQ	Fixed-Order Quantity
FOP	Fixed-Order Period
GA	Genetic Algorithm
GP	Gaussian Process
HSS	Hard Safety Stock
IPA	Infinitesimal Perturbation Analysis
KDE	Kernel Density Estimation
KPI	Key Performance Indicator
LR	Likelihood Ratio
LS	Lot-Sizing
LT	Lead Time
MCS	Monte Carlo Simulation
MDP	Markov Decision Process
MILP	Mixed-Integer Linear Programming
MIP	Mixed-Integer Programming
MOQ	Minimum Order Quantity
MRP	Material Requirements Planning
MRP-Opt	Material Requirements Planning Optimisation tool
MPS	Master Production Schedule
OSI	Optimisation with Simulation-based Iterations
PDF	Probability Density Function
PG	Parameter Group
PI	Probability of Improvement
PSO	Particle Swarm Optimisation
RL	Reinforcement Learning
RSM	Response Surface Methodology
R&S	Ranking and Selection
SA	Simulated Annealing
SAA	Sample Average Approximation
SBO	Simulation-Based Optimisation
SBM	Simulation Budget Management
SDA	Sequential Decision Analytics
SGD	Stochastic Gradient Descent
Sim-Opt	Simulation-based Optimisation framework

Abbreviation	Definition
SS	Safety Stock
S&OP	Sales and Operations Planning
TS	Tabu Search
UCB	Upper Confidence Bound
VNS	Variable Neighbourhood Search

1 Introduction

This chapter provides context about the topics discussed in this research. It will state the problem on which this study is based and defines the research scope. After defining the research questions, the use case on which the research will be applied is explained. Finally, the proposed research methods are presented.

1.1. Context

An important process in production systems is inventory control, as a flawed replenishment policy can result in material shortages or unnecessarily high stock levels [2]. An approach to inventory control in supply chains is Material Requirements Planning (MRP). MRP is defined as an intelligent information and decision-making system. It is designed to determine quantities and procurement timing of raw materials needed in a manufacturing system's production schedule. Given that material costs frequently constitute 60% to 90% of total production costs [3], the effective management of inventory through MRP is essential for profitability. MRP is an inventory control system which converts a Master Production Schedule (MPS) into material requirements and planned order schedules for inventory and production control [3]. To do this, it uses information on inventory, lead times, lot-sizing and the bill of materials (BOM). A BOM is a document that describes all the materials, components and instructions for manufacturing or repairing a specific product [4]. The demand for components is derived from the production levels of the end product, making MRP a dependent demand planning tool.

Three levels of planning exist in supply chain management. The strategic level represents long term planning under a high degree of uncertainty, considering all processes in an organisation. A level below that is tactical planning, concerning medium term planning where some more certainty exists. At the lowest level, operational planning considers day to day plans which are very detailed. At this level, MRP typically assumes a static product structure. The BOM and other information are treated as given inputs while the system computes the requirements and proposes the order schedules. This system works well for products with a stable design. However, many fast-moving consumer goods (FMCG) have frequent changes in packaging and product selection. This creates the need for changeover materials in the supply chain, which are needed to transition a production line from producing one item to the next. Packaging lines, for example, have to perform material changeovers for new packaging designs or specific packaging and labels for marketing campaigns. Non-changeover materials on the other hand, are standard, predictable materials which are consistently used in the operation. Producing FMCG can therefore make assuming a static product structure at the operational level invalid.

At the tactical planning level, MRP tools must therefore incorporate material changeovers. These transitions, sometimes described as rollover planning, must be timed very accurately. Timed too early, previously used materials may become obsolete. Timed too late, and the marketing campaign misses the intended window [5]. Ideally, the new product would be introduced as the old inventory depletes, since leftover inventory results in write-offs. Reaching this ideal is difficult because of uncertainty in demand and supply [6]. It requires strict coordination between marketing and operations departments, as marketing calendars set hard constraints, while operations try to execute a feasible and cost-efficient production. Processes which try to improve the coordination, such as Sales and Operations Planning (S&OP), also require quantitative decision-making when a supply chain becomes too complex.

An advanced way of formulating MRP problems is through Mixed-Integer Programming (MIP). MIP is used to solve problems in which both discrete and continuous decision-making occurs. An example of a discrete decision is whether or not to perform a changeover to a new material in a given week, while the number of produced pieces is typically modelled as a continuous variable. A MIP model tries to minimise or maximise a certain objective function, such as costs. MIP-based production planning is a very common modelling method. In practice, the MIP models are frequently run in a sequential, rolling horizon manner, for example as weekly re-optimisations. This way, updated information on demand, inventory positions and revised material requirements can be incorporated into the model over time. Although a MIP model is solved exactly, the decision-making environment which it tries to represent is not deterministic. Changes in demand, lead time uncertainty and disruptions in the supply chain make planning in these environments stochastic. Future inventory levels are for example affected by the present uncertainty and the decisions made today. A common way to work around uncertainty is by adding robustness to the MIP model. This can be done with safety stocks, or buffer inventory, which reduce the risk of running out of materials under uncertainty [7]. The buffers provide more of an indication of uncertainty than a complete representation of it. Parametrising safety stock levels presents a trade-off between safety stock costs, obsolete material costs and rush order costs. While having having higher safety stock levels causes more inventory costs, it could prevent rush orders in certain situations, preventing costs related to last-minute orders. At the same time, these larger inventories can also cause material obsolescence for changeover materials, resulting in costs for disposal for example.

Historically, planning parameters are often been calibrated only once, leading to performance degradation as environmental conditions evolve and variations occur. This leads to viewing the weekly re-optimisation process not as an isolated deterministic problem, but as an approximation of a Sequential Decision Analytics (SDA) problem. SDA problems try to find a policy which performs well over time and under uncertainty [8]. They do not necessarily provide a strong optimum for a single run. SDA and modern stochastic optimisation frameworks highlight that the performance of a problem should be assessed across sequences of decisions. This can be done using simulation and designs that explicitly model uncertainty, rather than treating it only through buffers. Thus, the research focus shifts from optimising one deterministic solution to re-optimising a decision process over time under uncertainty. This is especially relevant in a fast-moving production environment, where timing and material obsolescence risks are essential.

1.2. Problem statement

This research will study the problem of dealing with uncertainties in a stochastic production environment. MRP is widely used in manufacturing environments to translate production plans into material procurement and inventory decisions. While deterministic MRP models perform well for stable products, their effectiveness deteriorates in environments characterised by frequent material changeovers and stochastic demand. Therefore, a current shortcoming in MRP is the planning of changeover material stocks. Because of uncertainty in demand, costly rush orders which have minimum order quantities (MOQ) could be triggered. On the other hand, the stochastic demand can also cause material obsolescence. Despite these risks, industrial MRP tools often reduce safety stocks for changeover materials to zero in order to minimise obsolete inventory. The lack of safety stock for these materials increases the risk of running out of materials. This could already happen when demand patterns increase slightly compared to the current predictions. This problem is particularly concerning as this can lead to disproportionately high costs due to the need for last-minute ordering.

Advanced MRP tools increasingly rely on deterministic Mixed-Integer Programming (MIP) models with rolling horizon re-optimisation. However, the robustness to uncertainty is still largely addressed through static planning parameters such as safety stocks or lead times. Previous research shows that such parameters strongly influence system performance over time. It is therefore beneficial to determine these parameters from a sequential decision perspective rather than a deterministic setting. Most existing studies on this topic focus on standard materials and do not explicitly account for the unique risks associated with changeover materials. This provides a research gap, which this study tries to fill. This research addresses how safety stock levels for changeover materials should be determined in a stochastic MRP environment. The objective is to improve rolling horizon performance by developing a framework that calibrates safety stocks, while treating the existing deterministic MRP optimisation as a black box. To keep this research feasible, it will focus on creating a scalable framework, while testing with a small subset of materials and changeovers.

1.3. Research questions

To perform the study on the problem discussed in Section 1.2, the following main research question was established:

How can simulation-based optimisation be used to parametrise safety stocks of changeover materials in a deterministic MIP model, such that rolling horizon performance improves in a stochastic MRP environment?

In order to provide a thorough answer to the main research question, several underlying research questions will be answered throughout this report. How these questions will be answered, is discussed in Section 1.5. These questions will help with defining a conclusion, and are stated as follows:

RQ1. What is the state of the art on determining planning parameters in MRP under uncertainty?

- a) How is uncertainty in MRP currently modelled and mitigated in literature?
- b) How are safety stocks of changeover materials determined in literature?
- c) What planning parameters influence the robustness of the MRP environment?
- d) What methods are used to parametrise planning parameters in stochastic MRP environments?

RQ2. What is the current state of the manufacturer's MRP process?

- a) What are the key sources of uncertainty in the MRP changeover process?
- b) What do the uncertainties in the system look like?
- c) How are planning parameters currently determined?
- d) How are material changeovers modelled in the MRP-Opt tool?
- e) What are the objectives, constraints and decision variables?
- f) What are the determined safety stock levels of changeover materials?

RQ3. How can simulation-based optimisation be used to optimise the safety stock levels?

- a) What method(s) can be used to optimise the safety stock levels?
- b) What are the optimisation objective, variable and parameter(s)?
- c) What sources of uncertainty should be considered when parametrising the safety stocks?
- d) How can the discovered stochasticity be represented in the simulations?
- e) How do the tuned parameters compare to the currently used planning parameters (in terms of KPIs, computational performance, etc.)?
- f) How does the approach perform in other scenarios, such as different levels of uncertainty?

1.4. Use case

This study will be applied on the MRP of a large manufacturing company. The manufacturer operates several stochastic production environments throughout a global supply chain. It follows a B2B strategy, producing consumer goods made from raw materials, and selling them to wholesalers and retailers. Included in the supply chain are production, packaging, storage and distribution across multiple operating companies and facilities worldwide. A decentralised organisational structure plays a role in the complexity of planning within this manufacturing environment. A large product and brand variety, together with deviating production requirements and marketing campaigns per operating company, makes planning even more difficult. The planning environment relies on consistent decision-making and standardised data to reach desired performance and to minimise costs.

To tackle the complexity of planning changeovers, the manufacturer has invested in advanced planning and optimisation software. It partnered with specialised analytics providers, like ORTEC, to improve global planning performance in its supply chain. ORTEC developed an advanced tool for the manufacturer's MRP process specifically. The Material Requirements Planning optimisation (MRP-Opt) tool optimises material planning and changeover scheduling across a planning horizon, aiming to balance marketing constraints and material stocks across the full supply chain. It is built using Mixed-Integer Programming (MIP), which is a mathematical optimisation technique used to find the best solution to complex problems. The optimisation problem is expressed as a deterministic MIP model, with safety stocks used as a heuristic for robustness. The tool has provided substantial benefits in terms of planning efficiency and coordination. In this research, the MRP-Opt MIP is treated as a black box. This is due to the fact that it is not possible to analytically represent the MIP as the logic in this tool is very complex.

The manufacturer has acknowledged that the tactical MRP process could be more robust, especially to be able to handle demand fluctuations better and to mitigate the risks associated with the current safety stock approach. However, any enhancement of the tool's robustness must be done while keeping the optimisation tractable. In other words, the tool must maintain its efficiency and the ability to process large-scale supply chain decisions. Therefore, the goal is to develop a more adaptive safety stock approach that balances the trade-off between stock costs and ensuring operational robustness under uncertainty. Improving the robustness of the MRP process could help the manufacturer increase operational efficiency and reduce operational costs and the likelihood of disruptions in the supply chain.

1.5. Research approach

To answer the research questions in Section 1.3, several methods are proposed. RQ1 focuses on the state of the art on MRP under uncertainty and simulation-optimisation (Sim-Opt) techniques. An extensive literature review is performed to see which methods are used in literature and which research gaps the current literature has. Relevant sources are found using search queries and are sourced from scientific databases such as Scopus and ScienceDirect. The literature review is presented in Chapter 2.

RQ2 focuses on the current state of the manufacturer's MRP, and the material changeovers that happen within it. To analyse how the changeover process is currently planned by the manufacturer and modelled in its MRP-Opt tool, a current state analysis is performed. This analysis provides insights on how the safety stock levels are currently determined and how they can be improved. A domain model is developed which represents the connections between different components of the changeover process. The building blocks of the MIP underlying the MRP-Opt tool is described to the reader through a mathematical notation, including its objectives and relevant constraints. The current state analysis can be found in Chapter 3. To get an idea the uncertainties in the manufacturer's MRP system, a data analysis is performed. The data analysis aims to find the sources of uncertainty. It also aims to represent the stochasticity as realistically as possible. To do this, historical production plans are analysed based on a block bootstrapping method. The findings of this analysis help with representing the uncertainty in the proposed framework. Similarly, a material analysis is performed on the materials in the manufacturer's MRP, to be able to classify them according to their characteristics. The data analysis can be found in Chapter 4.

To answer RQ3, a more detailed explanation of the proposed methods is required, as this is the main body of this research. The aim of this research question is to find safety stock levels through the Sim-Opt procedure, testing different optimisation methods, assignments of uncertainty and settings. Section 1.6 illustrates the proposed Sim-Opt framework. The framework is explained through a conceptual model, based on which the actual model is built and explained in Chapter 5. Several experiments are performed in Chapter 6 to find reliable model settings and to test it in different scenarios. The results of these experiments are then discussed in Chapter 7.

1.6. Core methodology: Sim-Opt framework

To substantiate the decisions made in this research regarding the proposed Sim-Opt framework, the taxonomy proposed by Figuera and Almada-Lobo was consulted [9]. The taxonomy categorises Sim-Opt into purpose and hierarchical structure, among others. The simulation purpose of this study is to evaluate the solutions of the system that it represents. The system being the manufacturer's material changeover process within its MRP, while the solutions that will be evaluated will say something about the performance of this process. Section 2.2.5 will dive deeper into the different purposes for using Sim-Opt and how the two components can interact with each other. Section 2.2.6 describes a multitude of Sim-Opt techniques, all with their own advantages and challenges. The optimisation objective and decision variable, together with a more detailed explanation of how the Sim-Opt procedure works, is given in Chapter 5.

Regarding the hierarchical structure, the proposed Sim-Opt framework can be described as Optimisation with Simulation-based Iterations (OSI). This means that the simulation module is embedded within the iterations of an optimisation procedure. It can be used to evaluate the results of specific input settings. OSI is particularly suitable for complex systems characterised by stochastic behaviour and non-linear relationships. While OSI can offer a high level of modelling realism, it can be computationally demanding due to the repeated execution of simulation runs throughout the optimisation process. The effectiveness of OSI methods depends on efficient search strategies, careful management of simulation budgets, and the use of sampling techniques which balance solution quality and computational efforts.

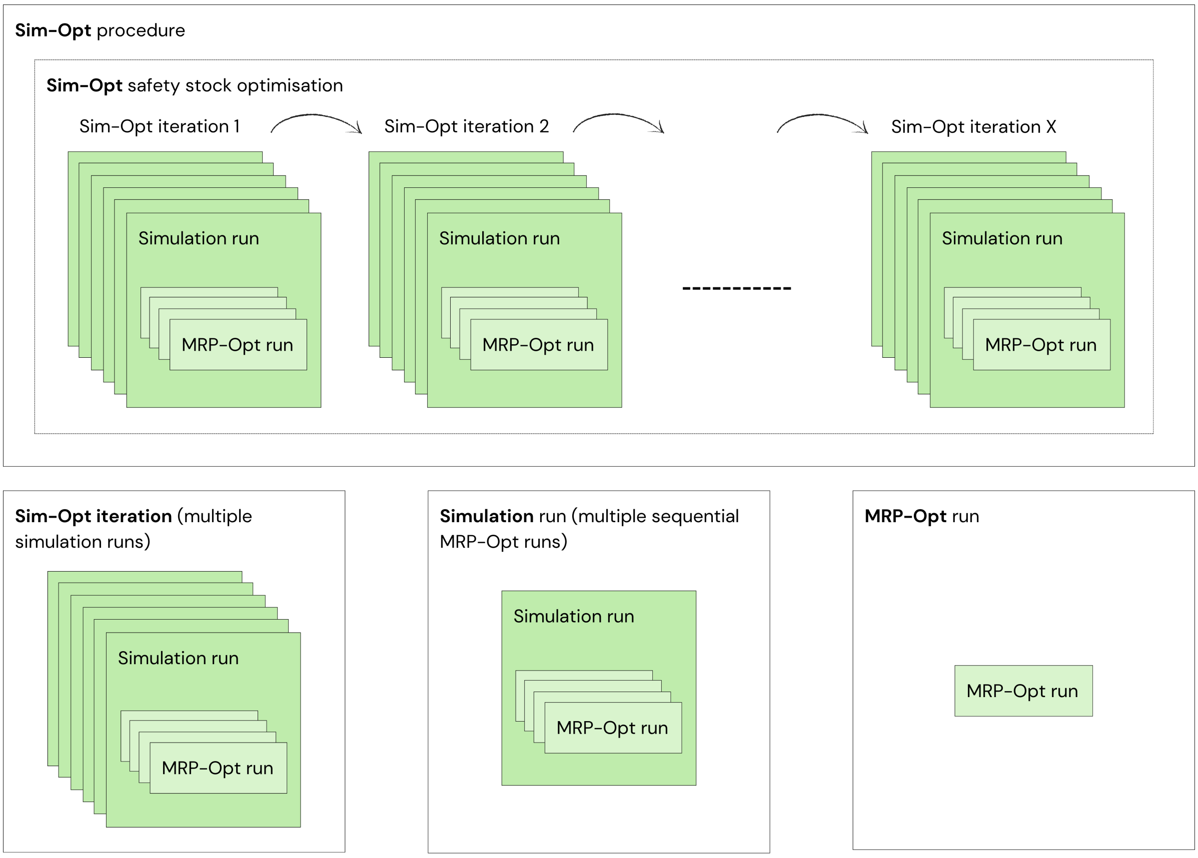

Figure 1.1: Sim-Opt framework

As can be seen in Figure 1.1, the iterative optimisation procedure is the leading component of the Sim-Opt framework. During each iteration, multiple complete simulations are run to evaluate solutions. Multiple simulations are run to cover the stochasticity and provide a reliable result. The simulation module serves as a performance evaluation mechanism in which the MRP-Opt model is run multiple times, imitating the Sequential Decision Analytics MRP problem at the manufacturer. The Sim-Opt objective is retrieved from this simulation module. Each MRP-Opt run represents one week, while a simulation represents the actual sequential MRP process that takes place at the manufacturer in a weekly manner. Figure 1.1 shows the hierarchy between the safety stock optimisation procedure, the simulations which represent the full MRP changeover process, and the executions of the MRP-Opt model which represent one instance of the sequential decision analytics (SDA) process.

1.7. Outline

The remainder of this thesis is structured as follows. Chapter 2 presents a literature review on Material Requirements Planning under uncertainty, safety stock parametrisation, and simulation-based optimisation techniques. Chapter 3 analyses the current state of the manufacturer's MRP process and describes the MRP-Opt model used in practice. In chapter 4, a data analysis of production plans and materials is provided, focusing on uncertainty modelling and material classification. Chapter 5 introduces the proposed simulation-based optimisation framework and details the modelling choices made in this research. The experimental setup is presented and the obtained results are discussed in chapter 6. Chapter 7 reflects on these findings, discusses their implications and relates them to the existing literature and practical application. Finally, Chapter 8 presents the conclusions of this research and provides recommendations for future work.

2
Literature review

This chapter will explain the search selection that was used to find relevant literature for this literature review. After this, it will analyse the state of the art on the parametrisation of MRP systems under uncertainty and the different techniques for doing this. Finally, it will present the contributions of this study to contemporary literature.

2.1. Literature search selection

To find relevant sources for this research, the search selection in Table 2.1 was used. The Table describes the concept groups for which keywords were selected. The truncation shows how the concept groups were used to find relevant sources in search engines. Most sources were found on Scopus and ScienceDirect, while others were found using forward and backward snowballing. No additional selection criteria were consistently used.

Table 2.1: Conceptual and methodological framework for literature review

Concept groups	Material Requirements Planning; SDA methods; Uncertainty in supply chains; Simulation-based optimisation
Keywords	**Material Requirements Planning:** MRP; inventory; inventory management; inventory control; tactical; material changeover; minimum order quantity; MOQ; supply chain; planning; changeover; network flow **SDA methods:** sequential decision analytics; SDA; mixed-integer programming; MIP; mixed-integer linear programming; MILP; deterministic; stochastic; parameter; parameter tuning; parametrisation; reparametrisation; data-driven; learning; machine-learning; machine learning; reinforcement learning; Markov Decision Process; MDP **Uncertainty in supply chains:** redundancy; safety stock; stock; buffer inventory; buffers; uncertainty; robust; robustness; demand fluctuations **Simulation-based optimisation:** simulation; simulation environment; simulation-based; optimisation; technique; method; algorithm; gradient-based; surrogate-based; metaheuristic; simheuristic; statistical; hybrid approach
Truncation	(Material Requirements Planning) AND (SDA methods) (Material Requirements Planning) AND (Uncertainty in supply chains) (Material Requirements Planning) AND (Simulation-based optimisation) (SDA methods) AND (Uncertainty in supply chains) (SDA methods) AND (Simulation-based optimisation (Uncertainty in supply chains) AND (Simulation-based optimisation)

2.2. Literature review

2.2.1. Inventory control systems

Besides MRP, several other production planning and control systems exist. These have been extensively analysed in literature, comparing the performance and resilience of the different strategies. Jodlbauer and Huber (2008) researched the robustness and stability of MRP, Kanban, CONWIP and DBR [10]. To assess robustness, they looked at how the systems performed in dynamic environments, while for stability they checked the sensitivity to parameter changes. Their findings show that CONWIP often achieves the highest service levels under optimal conditions. However, it lacks robustness when uncertainty increases in the environment it operates in. MRP demonstrated the highest stability of all the strategies, which means its performance is least sensitive to small changes in parameters. However, the research also noted that MRP, a push-based system, can be less robust than other pull-based systems. This happens mostly when MRP faces immediate variations, as the system only accounts for these changes during the next scheduled planning run.

2.2.2. Material changeovers and lot-sizing

Material changeovers refer to the set of activities required to transition a production line from processing one material or product variant to another. This process is particularly relevant in manufacturing environments with high product variety or short production runs, where changeovers happen frequently [11]. The changeover process typically begins with the termination of the ongoing production run. Operators halt machines, secure equipment, and remove any remaining materials from the preceding production batch. This is done to prevent contamination and the waste of materials. It also makes sure that the quality of the next product is guaranteed. This stage often involves cleaning parts of the production line, such as conveyor belts and mixing vessels. Next, the process moves into the setup configuration phase. This phase makes sure that machines, tools, or processing parameters are adjusted to match the specifications of the new product and its materials. This may involve mechanical adjustments, such as replacing molds, as well as digital or parameter-driven changes, such as updating machine programs or modifying feed rates.

Once the mechanical and digital settings are up to date, operators often perform a trial run. During this calibration phase, a small quantity of the new product is processed to verify that the equipment operates according to required quality standards. Technicians or quality inspectors may measure output characteristics, which depend on the product. If deviations are detected, corrective adjustments are executed before full-scale production starts. This ensures that the production line is compliant with predefined requirements. The final stage of the changeover process consists of documentation and process confirmation. Operators record the completion of the changeover and communicate the updated production status in digital systems such as Enterprise Resource Planning (ERP) platforms. These platforms are used to enable real-time visibility of the production process.

Material changeovers, often referred to as product rollover strategies, are an important driver of planning complexity in already complex supply chains. They often result in material loss and unnecessary costs. Lim and Tang researched the optimal timing and pricing for single-product and dual-product rollovers [5]. With single-product rollovers, phasing-out the old product and introducing the new product happens simultaneously. Dual-product rollovers allow for temporary co-existence of the two products in the supply chain. According to their research, dual-product rollovers are optimal when the marginal costs of old and new products are similar.

Another relevant research area is about the differences between theoretical MRP models and operational constraints. Javadi et al. analysed different lot-sizing methods in previous literature [3]. With lot-sizing, the optimal quantity to produce in each time period is determined. They identified that traditional lot-sizing methods, such as Economic Order Quantity (EOQ), often fail in practical applications. This is because they do not account for constraints which are specific for suppliers, such as MOQs and packaging sizes of batches. Their research developed a novel lot-sizing method and production scheduling algorithm that integrates these constraints into the MRP system. This provides a more pragmatic solution that prevents fragile planning outputs, which could lead to stockouts or obsolescence of materials when demand is uncertain. Fiorotto et al. researched the impact of setup carryover and setup crossover in lot-sizing problems [12]. A setup carryover makes it possible to begin a period with the previous item's setup state, while in a setup crossover, a setup operation can be split across two consecutive periods. Their study demonstrated that allowing for flexible setup decisions results in more flexibility to find better solutions and reduced operational costs.

2.2.3. Sources of uncertainty in MRP

Uncertainty in supply chains has been studied extensively in literature. Tordecilla et al. performed a literature review on scientific studies about Sim-Opt methods for designing supply chains under uncertainty [13]. Their study shows that cost optimisation is often the objective, while multi-objective approaches remain rare. The review revealed that demand uncertainty is most frequently researched in supply chain network design, far more than other sources of uncertainty such as capacity and costs.

Different sources of uncertainty have been studied within MRP specifically as well. Demand uncertainty, for example, captures flawed demand forecasts and the variability in customer orders [14]. It could lead to material stockouts or obsolescence, especially considering the bullwhip effect. Supply lead times provide another uncertainty, which can result in delays and inventory shortages. Yield uncertainty stems from problems with capabilities or product quality [15]. While other sources of uncertainty in MRP exist, the three mentioned risks are most frequently considered in current research. To mitigate these uncertainties, buffers can be implemented in the supply chain. For demand and lead time uncertainty, safety stock and safety lead time are commonly used to provide a buffer. To mitigate yield uncertainty, supplier diversification could be an option [16].

For uncertainty in MRP, similar trends as revealed by Tordecilla et al. [13] are seen. Barros et al. performed a systematic literature review about safety stock dimensioning under uncertainty and risk in the procurement process [15]. It selected relevant publications and categorised them according to the researched sources of uncertainty. The review shows that demand uncertainty is considered mostly when dimensioning safety stocks in MRP environments. In contrast to lead time and yield uncertainty, which have not been studied extensively. Combinations of demand and lead time uncertainty, and demand and yield uncertainty were often considered in research on multiple uncertainties.

2.2.4. Parametrisation of MRP under uncertainty

This section focuses on determining planning parameters for MRP under uncertainty. Scientific publications will be analysed and the state of the art about methods used to solve MRP problems under uncertainty will be discussed. As literature reviews on this topic, such the study of Barros et al. [15], have been performed in the past, this literature review aims to find relevant studies which were mostly published recently. Table 2.2 presents recent literature together with the type of planning parameter analysed, the objective of the study and the applied methods.

Table 2.2: Recent studies on MRP under uncertainty ('X' when parameter is considered)

Study	**Planning Parameter**			**Objective**	**Method(s) applied**
	SS	**LT**	**LS**		
Seiringer et al. (2025) [17]	X	X	X	min. total costs	Simulation
Werth et al. (2022) [18]	X	X	X	multiple	Simulation-based optimisation & Metaheuristics
Seiringer et al. (2022) [19]	X	X	X	min. total costs	Simulation-based optimisation & Metaheuristics
Seiringer et al. (2021) [20]	X	X		min. total costs	Metaheuristic & simulation
Thevenin et al. (2021) [21]	X		X	min. total costs	Two-stage & Multi-stage stochastic optimisation
Gansterer et al. (2014) [22]	X	X	X	service- & inventory level	Simulation-based optimisation

Abbreviations: safety stock (SS), lead time (LT), lot-sizing (LS).

Determining planning parameters of MRP systems facing uncertainty is a critical challenge in production environments. The values of these parameters has immediate impact on service and inventory levels, and effects the total operational costs. Traditional MRP often assumes deterministic demand. This requires the selection of planning parameters to be independent to create buffers which cover stochastic variability. Recent literature highlights the importance of integrated approaches to replace manual parameter selection, especially in highly complex, multi-echelon supply chains. Thevenin et al. for example, suggests that integrating the computation of lot sizes and safety stocks in a single framework provides better results, compared to computing them separately [21]. To overcome the sub-optimality of determining these parameters separately, they proposed multiple stochastic optimisation models. Their study compared two-stage (static) and multi-stage (dynamic) models, proving that multi-stage programs provide better reactivity. Their findings confirm that integrating these computations into a single framework removes the imaginary gap between planning tasks. This results in substantial cost savings compared to traditional safety stock policies, such as the base-stock policy of Graves and Willems [23].

Gansterer et al. investigated the efficiency of various optimisation methods in a make-to-order environment [22]. They found that a Variable Neighborhood Search (VNS) outperformed other methods in effectively searching large parameter spaces. As the dimensions of production systems increase due to multiple materials, computational complexity becomes a substantial problem. Werth et al. tries to address this problem by utilising a metaheuristic algorithm for multi-objective optimisation, combined with surrogate-assisted filtering and parameter clustering [18]. Their research showed that grouping materials by technology or production volume allows for high-dimensional optimisation of a real world scenario. Their approach can outperform current planning strategies while maintaining computational tractability.

The dynamic nature of demand forecasts in a rolling horizon environment often causes system nervousness. Frequent updates in demand forecasts lead to unstable production plans over time. Seiringer et al. explored this by developing a safety stock exploitation heuristic, designed to mitigate the unfavourable effects of periodic forecast updates in MRP [17]. Their study found that the Fixed-Order-Quantity (FOQ) lot-sizing policy is generally superior to the Fixed-Order-Period (FOP) policy under high uncertainty. This is because a FOP policy triggers inefficient, small-volume orders. Their results also provide the interesting insight that underestimating demand is significantly more costly for a manufacturer than overestimation. This is due to the higher cost penalties for backorders carry, compared to costs for holding excess inventory. Research into the rolling horizon approach supports viewing MRP as a SDA problem. Campuzano-Bolarín et al. researched a rolling horizon simulation, together with lot-sizing approaches, for managing demand with lead time variability [24]. Their research used system dynamics simulation to update real-time demand information and re-plan orders at a minimum cost. The key outcome of this study is that periodic re-planning generates better results regarding total costs and the reduction of the bullwhip effect compared to non-sequential methods.

Recent studies have introduced simheuristics to help define MRP parameters with limited simulation budgets. Seiringer et al. integrated Simulated Annealing (SA) and a Simulation Budget Management (SBM) concept to enhance search robustness [19]. The SBM concept improves search efficiency by skipping non-promising simulation iterations, leaving computational capacity for exploring broader larger parameter spaces. The key outcome of this research is that the combination of SA and SBM identifies parameter configurations that most effectively cover uncertainty, resulting in the lowest overall costs. Another research, by Seiringer et al., explored computational experiments which reveal that the proposed advanced heuristic extensions, significantly outperform basic heuristic parameter selection [20]. Integrating the second extension indicated that the algorithm can rapidly navigate complex search spaces. It identifies near-optimal parameter configurations under a limited simulation budget. These results highlight the potential for simheuristics to mitigate the negative effects of demand and lead-time uncertainty in multi-stage production environments.

Several advanced methodologies have been researched to solve sequential planning problems under stochastic conditions. Table 2.2 shows that in recent literature, combinations of simulation-based optimisation (SBO), metaheuristics and multi-stage stochastic programming were used to solve problems in MRP environments. Another methodology, called reinforcement learning (RL), is a widely used method for capturing uncertainty in supply chain management [25]. RL is a data-driven approach in which an autonomous agent learns to make decisions in an unknown environment by reacting to the associated outcomes [26]. In a study by Wang et al., RL was combined with a multi-agent simulation model as a strategy for dynamic inventory replenishment in a complex manufacturing environment [27]. Their research demonstrated that an RL-based agent can learn to adjust inventory parameters dynamically. This resulted in significant improvements in supply chain performance compared to static policies.

2.2.5. Simulation-based optimisation

Within hybrid Sim-Opt frameworks, the specific purpose assigned to the simulation component can be seen as a distinction between various research areas. According to Figueira and Almada-Lobo, these purposes can be categorised into four distinct roles: evaluation function , surrogate model construction, analytical model enhancement, and solution generation [9]. The first two categories are described as solution evaluation approaches, which can be used when the underlying response surface of a problem is unknown. For these purposes, simulation is used iteratively to assess the performance of candidate solutions, thereby guiding search moves and validating optimisation progress. Surrogate models use simulation output to build a problem-independent approximation of the system that is researched. This is searched directly by the optimisation module to find favourable regions, without having to run the computationally demanding underlying problem. Analytical model enhancement is utilised when a problem-specific analytical model exists but requires refinement. Examples are adjusting capacity parameters based on waiting times for a capacity allocation problem or extending a model to incorporate stochastic scenarios. Finally, solution generation involves simulation as the primary vehicle for producing the final solution. Optimisation is then used for internal computations without the need for simulation feedback to select between alternatives.

Figueira and Almada-Lobo further refine the interaction between the simulation and optimisation modules by a hierarchical structure [9]. This structure defines the functional dependence and the frequency of feedback between the two components. Figure 1.1 shows the hierarchical structure plotted against the different simulation purposes. Optimisation with simulation-based iterations represents a configuration where complete simulation runs occur within the iterations of an optimisation procedure to evaluate the quality of specific solutions. The alternate Sim-Opt structure involves both modules running in an alternating fashion, where either or both may run completely or incompletely in each iteration to progressively refine the results. Sequential Sim-Opt follows a simple linear chronological order where one component executes once, and its output is used by the following component to conclude the process. Simulation with optimisation-based iterations is a structure where a complete optimisation procedure is performed within the broader execution of a simulation process. The optimisation procedure can for example perform specific sub-tasks. Moving through these four categories, a trend can be seen. While the optimisation procedure gets more prominent, the total frequency of simulation runs relative to optimisation iterations decreases.

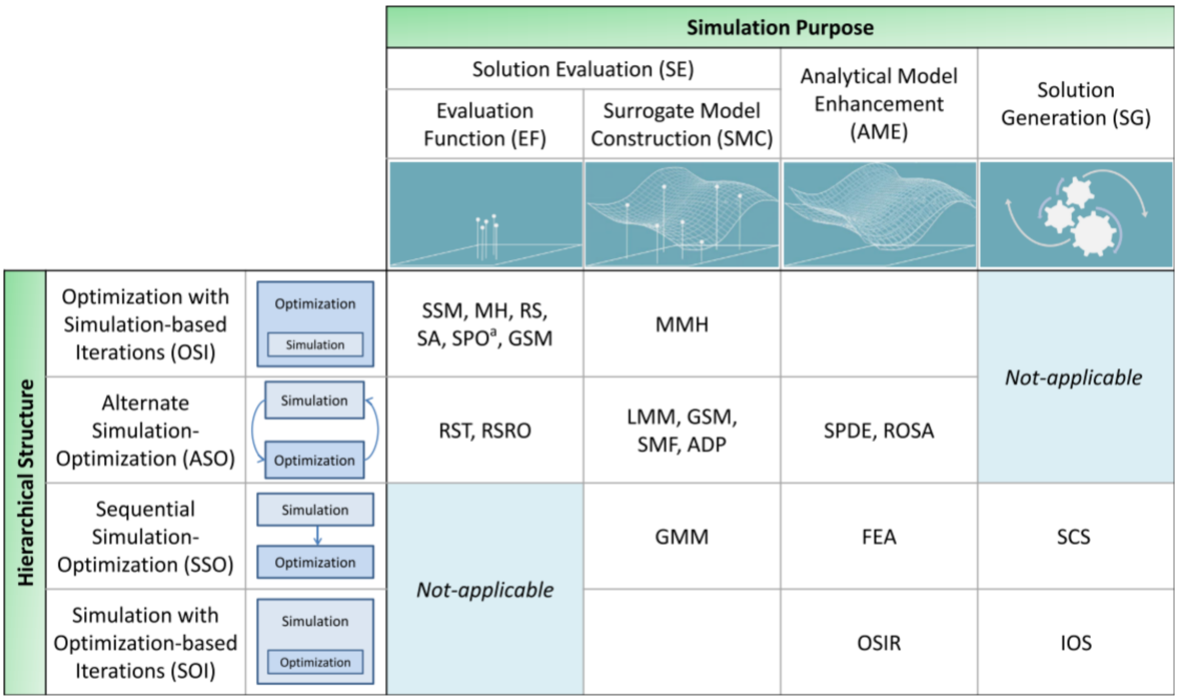

Figure 2.1: *Note.* From "Hybrid simulation–optimization methods: A taxonomy and discussion.", by G. Figueira, and B. Almada-Lobo, (2014), *Simulation Modelling Practice and Theory, 46,* p. 118–134 (https://doi.org/10.1016/j.simpat.2014.03.007)

2.2.6. Sim-Opt techniques

In current literature, many different simulation-based optimisation techniques have been studied. This study covers a wide range of different methods, but will not describe every available technique. While different ways of categorising these techniques exist, Table 2.3 shows how they are categorised in this study.

Table 2.3: Overview of Sim-Opt techniques

Method category	Techniques
Gradient-based	Simultaneous Perturbation Stochastic Approximation (SPSA), Infinitesimal Perturbation Analysis (IPA), Likelihood Ratio (LR)
Surrogate-based	Response Surface Methodology (RSM), Kriging, Bayesian Optimisation (BO), Artificial Neural Network (ANN)
Metaheuristics	Genetic Algorithms (GA), Particle Swarm Optimisation (PSO), Simulated Annealing (SA), Tabu Search, Variable Neighbourhood Search (VNS), Ant Colony Optimisation (ACO), Deep Q-Networks (DQN)
Statistical Methods	Ranking & Selection (R&S), Sample Average Approximation (SAA), Monte Carlo Simulation (MCS)

Gradient-based algorithms

Gradient-based algorithms are methods that iteratively improve decision variables by following estimates of the gradient of the objective function. This enables efficient navigation of high-dimensional continuous search spaces. Gradient-based algorithms are primarily utilised when the objective function is differentiable [28]. These algorithms operate through an iterative update mechanism, where an estimate of the objective function's gradient is used to determine the optimal search direction toward an optimal solution [29]. The general mathematical framework typically involves a stochastic approximation approach, which can be categorised into Robbins-Monro or Kiefer-Wolfowitz algorithms. The former when unbiased gradient estimators are available, the latter when estimators are only asymptotically unbiased, such as those derived from finite-difference approximations [29]. The iterative process adjusts decision variables by a step size proportional to the estimated gradient.

The effectiveness of these optimisation routines is highly linked to the precision of gradient estimation, frequently referred to as sensitivity analysis. Direct estimation techniques, such as infinitesimal perturbation analysis (IPA) and the likelihood ratio (LR) or score function methods, are widely utilised because they often yield unbiased estimators [28]. IPA functions by computing a sample-path gradient estimator simultaneously with the simulation process, making it an efficient choice for various stochastic models [29]. The LR method uses knowledge regarding the probability distributions of the underlying random variables and the sample performance functions to derive gradient settings [29].

For high-dimensional optimisation problems where direct methods may be computationally expensive or inapplicable, indirect techniques like simultaneous perturbation stochastic approximation (SPSA) offer a robust alternative. SPSA is distinguished by its extreme computational efficiency, as it requires only two objective function measurements per iteration to approximate the gradient regardless of the number of decision variables involved [30]. This feature provides a significant advantage over traditional finite-difference stochastic approximation, which requires a number of evaluations proportional to the problem's dimensionality. Furthermore, in the context of large-scale simulation optimisation, stochastic gradient descent (SGD) has emerged as a predominant method because its optimisation accuracy is independent of the total dataset size, making it computationally superior to standard gradient descent for managing massive datasets [29].

Surrogate-based algorithms

Surrogate-based algorithms, also referred to as metamodelling, are designed to address the challenges of computationally expensive simulation models. They do this by substituting the simulation model with a less exhaustive mathematical approximation. In simulation-based optimisation, the evaluation of a primary numerical model can require significant time and resources, which makes traditional iterative search methods impractical. Surrogate models provide an approximating function for the system's performance that is inexpensive to compute, facilitating a more efficient exploration of the decision space [28]. The use of these algorithms introduces a structural element of error that must be taken into account for accurate decision-making. However, they are essential for optimising complex real-world systems where the objective function is analytically intractable and direct simulation is expensive [31, 32].

Common methodologies within this field include Response Surface Methodology (RSM) and Kriging, which is also known as Gaussian Process Regression. RSM uses a set of statistical and mathematical techniques, typically involving local linear or second-order regression models, to identify optimal operating conditions. It does this by fitting models to specific regions of the search space [22]. Kriging serves as a global metamodel that approximates outputs across the entire search space through a predictor [31]. An important feature of Kriging is its ability to quantify prediction uncertainty through the mean square error, the so-called Kriging variance. This variance represents the likelihood of a response given the observed data. Using an infill criterion, Kriging allows algorithms to balance local exploitation and global exploration as they sequentially select new points to simulate and update the model [31].

Bayesian optimisation (BO) is a sequential strategy employed to optimise objective functions that are analytically intractable and computationally expensive to evaluate. This approach is particularly relevant for simulation-based optimisation (SO) problems where the objective or constraints lack an analytical form and must be estimated through simulations. The methodology is based on the construction of a surrogate model, often a Gaussian process, which serves as a computationally efficient approximation of the expensive simulation [29]. The ability to quantify uncertainty is an important characteristic of Bayesian optimisation, as it allows the algorithm to focus its efforts on areas where information is lacking [18]. Artificial Neural Networks (ANN) function as a supervised learning technique inspired by the human nervous system [33]. This method is capable of building non-linear relationships between decision variables and simulation results. ANN is widely recognised for its high accuracy and flexibility in treating problems with non-linear structures, although it generally requires larger datasets compared to other metamodelling approaches to ensure effective mapping and avoid overfitting [33].

Metaheuristics

Metaheuristics represent algorithms that are inspired by nature. They are designed to address complex real-world optimisation problems that are often analytically intractable. These algorithms serve as powerful tools for navigating decision spaces where the feasible set is either infinite or too large for a comprehensive simulation [28]. The effectiveness of a metaheuristic depends on the balance between local improvement and global exploration of the search space. By integrating metaheuristics in computer simulations, researchers can evaluate the performance of potential solutions under stochastic conditions without making the restrictive assumptions required for traditional mathematical models. Jalali and Van Nieuwenhuyse found that among simulation-based optimisation techniques, metaheuristics and hybrid techniques are most popular [28].

Population-based algorithms is one of the sub-categories in metaheuristics. They utilise multiple candidate solutions simultaneously to maintain diversity and to prevent premature convergence to local optima [34]. Genetic Algorithms (GA) are the most widely employed method in this group, which mimic the biological process of natural selection through iterative operations such as selection, crossover, and mutation [34]. Other population-based methods include Particle Swarm Optimisation (PSO) and Ant Colony Optimisation (ACO). PSO is inspired by the social behaviour of bird flocking or fish schooling, while ACO is modelled after the foraging behaviour of real ant colonies. In contrast to population-based approaches, single-solution metaheuristics focus on improving a single candidate solution through local search procedures. Simulated Annealing (SA) is a well-known method in this category, inspired by thermodynamics. SA avoids being trapped in local minima by occasionally accepting non-improving solutions based on a decreasing probability [35]. Tabu Search (TS) is another local search metaheuristic that uses memory mechanisms to discourage the search from revisiting recent solutions or becoming trapped in cycles. Furthermore, Variable Neighbourhood Search (VNS) operates by systematically changing neighbourhood structures during the search process to also escape local optima [22].

Recent literature has increasingly applied reinforcement learning (RL) techniques to inventory problems. According to research by Zabraoui et al., RL is found to be very adaptable and is able to handle dynamic uncertainties [36]. Deep Q-Networks (DQN) model inventory and resource control as Markov Decision Processes (MDP) to learn policies that maximise cumulative rewards [36]. To further enhance performance, these methods are frequently combined into hybrid frameworks, sometimes referred to as simheuristics. These hybrid methods integrate metaheuristics like GA or SA with adaptive learning and possibly intelligent simulation budget management to optimise systems in supply chain analytics [19].

Statistical methods

Monte Carlo Simulation (MCS) is a probabilistic sampling technique used within simulation-based optimisation, particularly useful for modelling non-dynamic systems [37]. This methodology can be combined with global search algorithms to estimate the relationship between input parameters and objective outcomes through iterative exploration [38]. Similarly, the Sample Average Approximation (SAA) method uses sampling techniques to replace an analytically intractable stochastic optimisation problem with a deterministic approximation. Ranking and Selection (R&S) procedures are designed for optimisation problems where the feasible set consists of a discrete and relatively small number of alternatives. The primary objective of R&S is to identify the single best solution with a pre-specified level of statistical confidence.

2.3. Contributions to literature

Table 2.4 describes the strengths and weaknesses of the literature discussed in Section 2.2.4. These studies are, to the best of the author's knowledge, the only studies which have recently researched the parametrisation of planning parameters in MRP under uncertainty specifically. The following paragraphs report the specific research gaps identified within the analysed literature that emphasise the relevance of the proposed research.

Table 2.4: Research gaps of analysed literature

Study	Strengths	Limitations & research gaps
Seiringer et al. (2025) [17]	Developed a safety stock exploitation heuristic, which can mitigate the negative impact of incorrect demand predictions;	Future work could explore more advanced optimisation algorithms and a more complex application
Werth et al. (2022) [18]	Application on a high-dimensional real world case; Multi-objective approach	Future research could address other objectives and constraints
Seiringer et al. (2022) [19]	Simheuristics used for parameter optimisation; Proposed an intelligent simulation budget management to skip irrelevant solutions	Does not consider interdependency of optimised planning parameters; Application is a simple production system; Used a single objective
Seiringer et al. (2021) [20]	Analysed a rolling horizon MRP with random demand, lead time and setup times	The use of a basic production system as application; Ignores the lot-sizing parameter
Thevenin et al. (2021) [21]	Joint optimisation of lot-sizes and safety stock; Compared a two-stage and multi-stage stochastic approach for different decision frameworks	Multi-stage model requires heavy computational time, research on stage-wise decomposition methods could address this problem; Ignores the lead time parameter
Gansterer et al. (2014) [22]	Used hierarchical planning; Evaluated/compared different optimisation methods	Influenced by typical make-to-order situations specific to suppliers in the automotive industry; Used identical parameter settings for all products

An important research gap in current production planning literature is the limited focus on specific types of uncertainty. The complexities associated with demand timing are often ignored. Recent simulation studies, such as Seiringer et al. [17], explicitly state that their models address demand quantity uncertainty while excluding demand shifting or changes in demand due dates. Similarly, earlier simheuristic frameworks acknowledge that other sources of uncertainty, besides order sizes and machine setup times, are frequently ignored in order to keep the model simple [20]. The timing of rollover windows and material changeovers are critical to minimise the risks of stockouts and material obsolescence.

Current research mostly focuses on the optimisation of safety stocks and lead times for standard, stable materials. Although studies like Lim and Tang [5] have researched optimal product rollover strategies, there is a notable lack of literature addressing the unique requirements of changeover materials in combination with parametrisation of other planning parameters. In existing industrial tools, safety stocks for these transitional materials are often reduced to zero to mitigate the risk of leftover inventory and resulting cost write-offs. This approach can be very efficient in deterministic scenarios. Nevertheless, this policy introduces extreme vulnerability to a supply chain in a stochastic environment. Small fluctuations in demand can trigger costly MOQs and high costs for last minute orders. This identifies a need for a comprehensive safety stock system that provides robustness during material transitions.

This study focuses on determining safety stocks for changeover materials in MRP under uncertainty. The manufacturer's MRP will be considered as an SDA problem, where planning parameters will be re-optimised in a sequential manner. There are multiple ways to solve SDA problems, including multi-stage stochastic programming and RL. Although efficient in certain environments, those methods don't scale well to industry-sized problems. As seen in the analysed literature, a promising approach is a combination of simulation-based optimisation and metaheuristics to determine planning parameters. In this study, the planning parameters will be optimised using Simulated Annealing in a stochastic simulation environment. The manufacturer's MRP provides a complex application for parametrising MRP under uncertainty. The combined approach of optimisation, simulation and other techniques, applied on a large-scale real world industry use case, could enhance current literature.

In this study, the MRP-Opt model used to optimise the manufacturer's MRP process, can be considered a black box. This is because the MRP-Opt model is very complex and it is not possible to analytically represent it. This also provides a research gap for this research, as the system which is represented by the simulation in other studies is often largely adapted. To the best of the author's knowledge, considering the MRP system as a black box in a Sim-Opt study, applied on an industry-size problem, is not common in literature.

3 Current state of the manufacturer's MRP process

This chapter describes the current state of the manufacturer's MRP process. It will discuss the demand uncertainty that plays a role in this process, and ways of mitigating this uncertainty. Next, it will elaborate on how the material changeover process is currently modelled by the manufacturer.

3.1. Manufacturer's MRP process

The MRP process at the manufacturer can be described as a sequential decision analytics (SDA) process. The manufacturer uses a MIP model in its MRP-Opt tool to optimise its MRP, described in detail in Section 3.2. The MIP is solved every week. New data is constantly fed to the MRP-Opt tool through constant daily or weekly data flows or manually by users at the manufacturer. Each week, the MIP is solved at the same moment, with updated data, and the planning cycle repeats itself the next week. At the end of each week, or planning cycle, a snapshot of the data at that moment is made and saved into a **casefile** for future reference. The model is used to make decisions over a longer time period. The MRP-Opt tool optimises over a horizon of 52 weeks, so 1 year into the future. The decisions made and actions taken influence the future state of the MRP process. The manufacturer's MRP process is used to look well ahead into the future, determining stock levels and procurement plans, and optimising changeover timing. They do this based on forecasted demands. This means that determined stock levels and procurement plans for a certain week intuitively get more accurate as this week is getting closer. Figure 3.1 describes the manufacturer's SDA process, where the MRP-Opt's MIP is run each week with updated information. With this data, the model optimises the changeover timing over a set period of time, 15 weeks in this example.

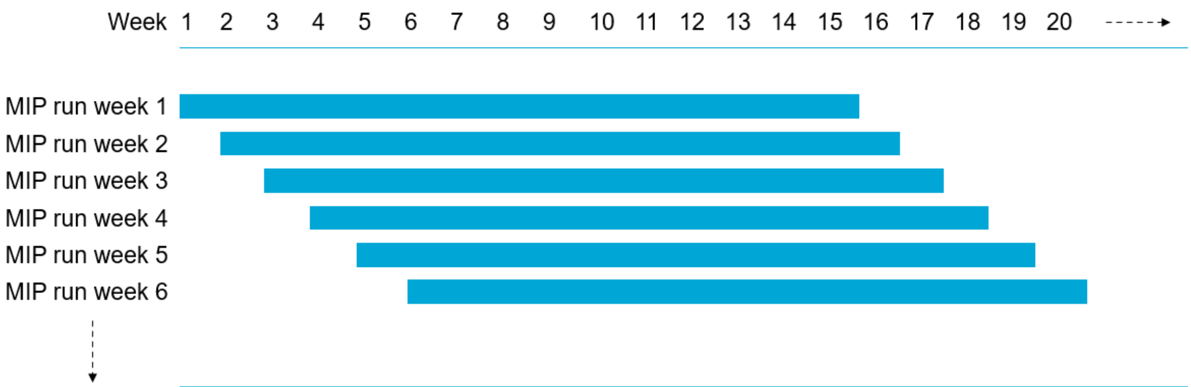

Figure 3.1: SDA process of the manufacturer's MRP

Figure 3.2 describes the different domains within the MRP process, and how they interact. A production plant receives a production plan, which indicates what quantity of a certain product or products the brewery needs to produce in a certain week. It also states on what resource in the production plant this product needs to be made. The production process, and implicitly the production plan, results in finished products. In return, the demand for finished products is input for future production plans. The production plan uses the bill of materials (BOM), which indicates what raw materials are needed and what the production process of a finished product looks like. The BOM determines how much of certain materials is used during production. If the production plan indicates that a product needs a new label, while last week it had another label, the production process requires a material changeover. Material usage depletes the stock levels, which can be located at the production plants, but also at vendors. The stock levels at all locations are taken into account during the MRP process. The material stocks can have different requirements, depending on their location. For example, the safety stock of a raw material can deviate per location. Depletion of material stocks trigger call-offs, which is a request for the delivery of a specific material, based on a predetermined contract. Having materials in stock costs money, as well as ordering materials at vendors and being left with obsolete stock. The costs which are a result of the MRP process, consist of these three components. Costs are determined weekly per material-location combination.

The MRP process works towards several objectives. They are determined based on a certain priority level, meaning one can be more important to the manufacturer than the other. These objectives are also placed within the domain model in Figure 3.2. The legend numbers the objectives in order of priority, while they are placed between the domains. For example, objective 1 is placed within the costs domain. Material changeovers can result in obsolete materials which will not be used any more, resulting in unwanted costs. Objective 1 tries to minimise these obsolete materials, and therefore the obsolete stock costs. The objectives are explained in more detail in Section 3.2.1.

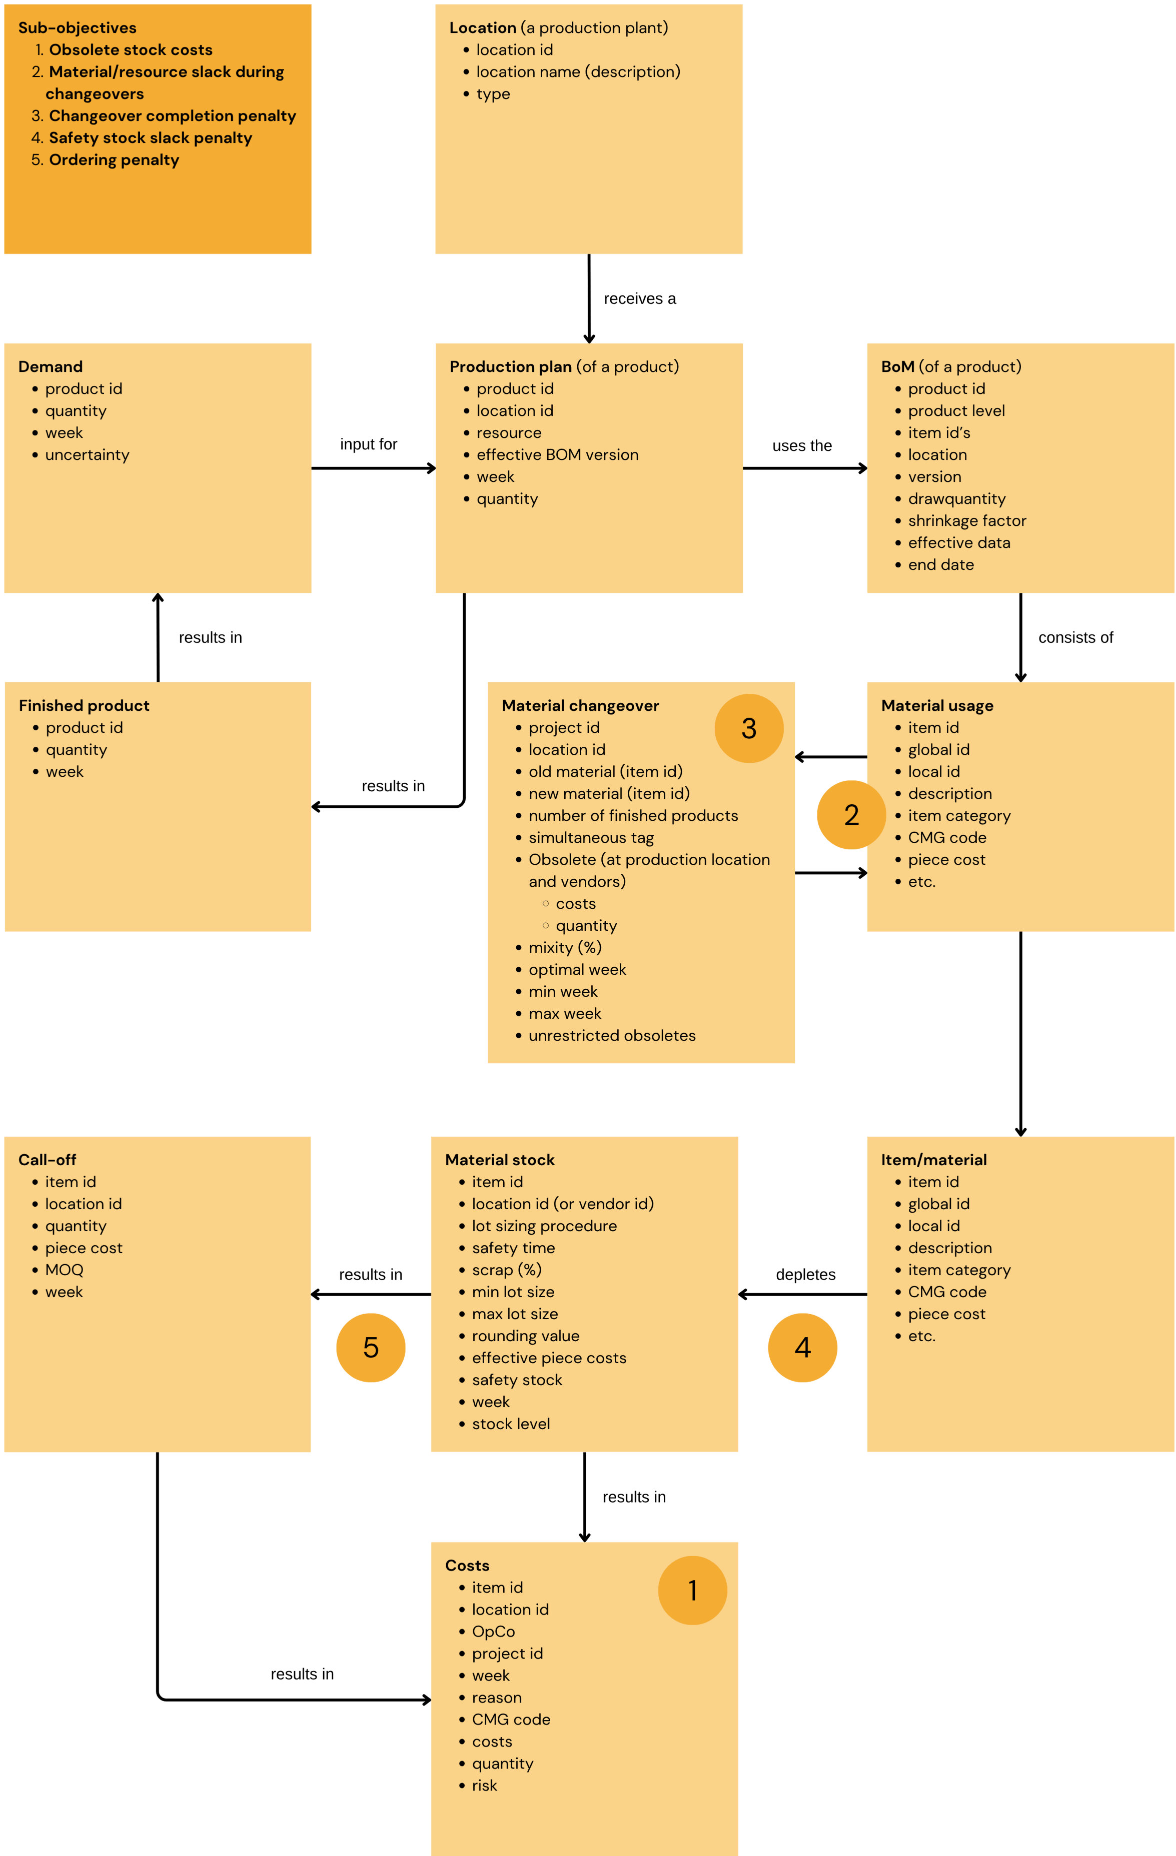

Figure 3.2: Domain model MRP at manufacturer (with 1-5 being the MRP-Opt sub-objectives)

3.1.1. Demand uncertainty

Demand uncertainty plays an important role in the stochasticity in Material Requirements Planning (MRP) environments. Especially in SDA settings, it can have a disproportionate impact on material planning decisions. It originates from inaccuracies in demand forecasts, variability in customer orders and changes in demand timing. In the manufacturer's MRP process, demand for materials is derived downstream from finished product forecasts through the bill of materials. This can cause forecast errors to propagate upstream and increase through the supply chain. This amplification is often called the bullwhip effect. It makes inventory levels and replenishment decisions already sensitive to small demand deviations. With a weekly re-optimised MRP model, such forecast updates can lead to changes in planned production and procurement volumes. This increases system nervousness and reduces the stability of material flows. As a result, demand uncertainty has a direct effect on inventory costs and the feasibility of planned production schedules.

3.1.2. Current safety stock policy

A planning parameter which has a great influence on MRP costs is safety stock. Safety stocks are a means of adding robustness in case of variant production demand. They can help reduce the risk of stockouts when a last minute spike in demand increases production levels. Stockouts slow down the production which costs money. On the other hand, they can also cause material redundancy. When there is no rise demand and the safety stock is therefore not used, safety stocks of phased out materials can become obsolete. Obsolete material stocks also come with a price, so there is a certain balance to be found concerning safety stocks. Meanwhile, having more materials in stock to act as a buffer also increases costs, as it is simply idle capital.

The current safety stock policy at the manufacturer tolerates the assigned safety stock levels to not be maintained in the MRP-Opt model. Because of the MRP-Opt objectives structure and safety stock constraint, explained further in Section 3.2.1, the model allows stock levels to dive below the predefined safety stock level. Therefore, the safety stocks currently maintained will be referred to as **soft safety stock**. The manufacturer currently determines safety stock values on a material-location level, meaning the safety stock of a certain material can differ per production location or warehouse. The safety stock levels have once been decided deterministically, meaning they are always the same. An important note is that currently not all materials have safety stocks assigned to them, or for example not at every location. Table 3.1 gives an idea of what the safety stock parameters look like, applied to a fictional example.

Table 3.1: Example safety stock levels

Material	Week 1	Week 2	Week 3	Week 4	.....
Packaging	1200	1200	1200	1200	.....
Labels	850	850	850	850	.....
Bubble wrap	0	0	0	0	.....
Tape	230	230	230	230	.....

3.2. MRP optimisation tool

To manage their MRP process, the manufacturer uses an optimisation tool, called MRP-Opt. This tool is responsible for planning material changeovers and assisting replenishment planners with their inventory management. The MRP-Opt model is built on a software platform on which it is possible to build and manage optimisation models. This platform makes it easy to build a user interface so that users can interact easily with the software and the optimisation model. The manufacturer's replenishment planners are therefore able to easily make changes to a production plan for example. The optimisation model is automatically solved weekly on a scheduled time, using the newly submitted input data. Manually triggered optimisations are possible any time as well and are mostly used by planners for testing different scenarios.

To optimise changeover timing, the tool considers all material stocks in the supply chain (at vendors and production locations) to identify the best time to switch from one material to another. It would obviously be most ideal if a changeover happens when the stock of the obsolete material is fully depleted. However, this is not always possible. Changeover materials can have a simultaneous tag, which represents a link with another changeover material. It requires the changeover of these materials to happen simultaneously. Therefore, it can happen that the stock levels of these materials are not fully depleted and the manufacturer is stuck with leftover stock. It is also possible to assign sequentiality to a set of changeovers. Sometimes it is necessary to let a production process switch one material first, after which another material can be switched later on. These changeovers do not have to take place immediately after each other, but the first one does need to happen before the other.

The material changeovers in the manufacturer's processes can have different structure types. The simplest structure is switching from material A to material B permanently, representing a shift to a new raw material in the production process. Due to temporary marketing campaigns, an A-B-A structure is also possible. Packaging is for example changed for a certain period, and later changed back to its original look.

3.2.1. MRP-Opt model

The MRP-Opt model optimises the changeover process using a deterministic mixed-integer programming (MIP) model. MIP modelling is a mathematical optimisation technique which concerns the systematic identification of the best possible solution to a mathematical problem within a feasible region. This region is defined by decision variables and a set of constraints. Some decision variables are restricted to integers, while others can be continuous. The performance is evaluated through an objective function that is either minimised or maximised. A basic optimisation problem can be expressed as:

$$\min f(x) \quad \text{subject to } g(x) \leq b, \tag{3.1}$$

where x denotes the decision variable, $f(x)$ the objective function, $g(x)$ a constraint function, and b a bound [39].

A MIP model can be described using a mathematical notation, which creates a simplified mathematical representation of a system. This notation, including sets, indices, parameters and decision variables, is described in Table 3.2.

Table 3.2: Mathematical notation of the MRP-Opt model

Sets and indices		
L	set of locations	$\ell \in L$
M	set of materials	$m \in M$
P	set of products	$p \in P$
R	set of production resources	$r \in R$
W	set of weeks	$w \in W$
C	set of changeovers	$c \in C$
D	set of valid changeover tuples satisfying time-window constraints	$(\ell, p, r, m_f, m_t, w) \in D$
Parameters		
p_w^{time}	ordering time penalty for week w	[–]
$p_{i,j}^{\mathrm{flow}}$	flow penalty for shipping from location i to j	[–]
$c_{\ell,m}$	effective material cost for material m at location ℓ	[€]
$u_{\ell,\,p,\,r,\,m_f,\,w}$	material usage at changeover c for tuple (ℓ, p, r, m_f, w)	[pieces]
Q_c^{fix}	fixed material usage for a changeover c	[pieces]
$SS_{\ell,m,w}$	safety stock level of material m at location ℓ in week w	[pieces]
$in_{\ell,m,w}$	incoming flow of material m at location ℓ in week w	[pieces]
$out_{\ell,m,w}$	outgoing flow of material m at location ℓ in week w	[pieces]
Decision variables		
$s_{\ell,m,w}$	stock level of material m at location ℓ in week w	[pieces]
$x_{i,j,m,w}$	flow of material m from location i to j in week w	[pieces]
$slack_{\ell,m,w}^{\mathrm{ss}}$	safety stock slack of m at manufacturing location $\ell \in L$ in week w	[pieces]
$c_{\ell,p,r,m_f,m_t,w}$	fraction of changeover completed for tuple $(\ell, p, r, m_f, m_t, w)$	[0,1]
$slack_c^{\mathrm{m}}$	material/resource slack for changeover c	[pieces]
$obs_{\ell,m,w}$	obsolete stock of material m at location ℓ in week w	[pieces]

The MRP model performs a multi-objective optimisation, consisting of the sub-objectives below. The sub-objectives are listed in order of priority. When the optimisation is run, it fully optimises the first objective, after which it continues with the next, and so on. This also means that when sub-objectives are contradictory, the objective with the highest priority is dominant.

1. **Obsolete stock costs:** This objective minimises obsolete stock costs by penalising excess materials that become obsolete, weighted by material cost.

$$\min \sum_{\ell \in L} \sum_{m \in M} c_{\ell,m} \cdot obs_{\ell,m,w} \tag{3.2}$$

2. **Material slack during changeovers:** This objective minimises the material slack variable by penalising violations of material usage constraints during changeovers.

$$\min \sum_{c \in C} slack_c^{\mathrm{m}} \tag{3.3}$$

3. **Changeover completion:** This objective minimises uncompleted changeovers by minimising the fraction of uncompleted changeovers within allowed time windows.

$$\min \sum_{(\ell,p,r,m_{\mathrm{f}},m_{\mathrm{t}},w) \in D} (1 - c_{\ell,p,r,m_{\mathrm{f}},m_{\mathrm{t}},w}) \tag{3.4}$$

4. **Safety stock slack penalty:** This objective minimises safety stock slack by penalising falling below safety stock at production plants.

$$\min \sum_{\ell \in L} \sum_{m \in M} \sum_{w \in W} slack_{\ell,m,w}^{\mathrm{ss}} \tag{3.5}$$

5. **Ordering penalty:** This objective minimises material flow by penalising placing orders too early and using costly transport routes.

$$\min \sum_{i \in L} \sum_{j \in L} \sum_{m \in M} \sum_{w \in W} p_w^{\text{time}} \cdot p_{i,j}^{\text{flow}} \cdot x_{i,j,m,w} \tag{3.6}$$

Together, these sub-objectives guide the model towards reducing obsolete inventory, avoiding costly slack in resource usage, efficient and timely changeovers, minimising stockouts and just-in-time ordering. The MRP tool optimises these objectives under a large number of constraints. Due to the many constraints, not all constraints can be described in this report. However, there are two constraints which are important to mention as they have a direct connection to the safety stock:

- **Safety stock constraint:** If a production facility ℓ uses a given material m in a given week w, then the available stock $s_{\ell,m,w}$ plus a slack variable $slack_{\ell,m,w}^{\text{ss}}$ must be at least the required safety stock level $SS_{\ell,m,w}$ of this material. This constraint can be seen as a soft constraint, as the model allows stock levels to fall below the established safety stock level value by adding the slack variable.

$$s_{\ell,m,w} \; + \; slack_{\ell,m,w}^{\text{ss}} \; \geq \; SS_{\ell,m,w}, \tag{3.7}$$

- **Location stock balance constraint:** The stock level $s_{\ell,m,w}$ of material m at location l in week w should be equal to the previous stock $s_{\ell,m,w-1}$ in week $w-1$ plus the incoming material flow $in_{\ell,m,w}$, minus the outgoing material flow $out_{\ell,m,w}$.

$$s_{\ell,m,w} \; = \; s_{\ell,m,w-1} \; + \; in_{\ell,m,w} - out_{\ell,m,w} \tag{3.8}$$

Because these constraints contain the same decision variable $s_{\ell,m,w}$, the MRP-Opt model has to find a value for this variable which satisfies both constraints. When a value is found for $s_{\ell,m,w}$ which does not satisfy the safety stock constraint, this is not a problem as the slack variable $slack_{\ell,m,w}^{\text{ss}}$ can become as large as necessary. As the safety stock constraint is not a hard constraint, the stock levels of certain changeover materials are sometimes optimised below the established safety stock level. The stock levels can even be reduced to zero. This happens due to the priorities of the sub-objectives. The decision variable $s_{\ell,m,w}$ is directly connected to the first sub-objective. The material stock of a changeover material that had its changeover, becomes obsolete material stock $obs_{\ell,cm,w}$. First and foremost, the optimisation tries to minimise obsolete stock costs, even if this means going below the safety stock level. Looking at sub-objective 4, the model also tries to minimise safety stock slack. As these are contradictory, but sub-objective 1 has a higher priority, the safety stock slack is taken for granted. When stock levels of these materials are optimised below the safety stock level, a risk arises. It can result in unsolicited call-offs, when the demand for this material suddenly increases above expectancy.

Some examples of other constraints present in the MRP-Opt model:

- **Location material usage balance:** this constraint makes sure that the materials used on a resource at a certain production plant is equal to the demand for this material according to the active BOM.
- **Changeover happens exactly once:** this constraint makes sure that a changeover can not happen more than once.
- **Simultaneous changeovers:** this constraint makes sure that changeovers which are paired through a simultaneous tag, actually take place at the same time.
- **Changeover sequentiality:** this constraint makes sure that a set of changeovers with sequentiality happen in the correct sequence.

The MRP-Opt model uses a lot of input data and also produces a lot of output data. Table B.2 in Appendix B mentions the output that the optimisation generates and what they mean. Table B.1 in Appendix B gives an overview of all the input data that the model requires.

4 Data analysis

This chapter will provide a data analysis of relevant data that is either necessary to build the Sim-Opt framework, or to use in the Sim-Opt procedure. It will first explain the data restrictions that this research had to work with. Then, the production plans will be analysed to capture the stochasticity in them.

4.1. Data availability

The manufacturer saves data in the form of casefiles, which include all available data for the MRP-Opt model to run. As described in Section 3.1, these casefiles are saved as snapshots of the data on a weekly basis. The manufacturer has been consistently saving casefiles on a weekly basis as of October 2025. Although the data of 2 weeks is missing, this creates a fairly consistent availability of data for about half a year. From before October 2025, very little data is available and is therefore disregarded. Not having historic data from at least one full year, and preferably more years, created restrictions for the data analysis. The reliability of the analysis is less solid because external influences and seasonal changes in demand fore example are not fully captured by the available data.

A material can be used at multiple locations, or used for multiple products at the same location. These material-location-product combinations are called **items**. These items could also be described as Stock Keeping Units (SKUs). Much of the data is available on an item level. While items and materials can sometimes by used simultaneously, there is a clear distinction.

4.2. Production plan analysis

The production plans were analysed to get a sense of the stochasticity involved. This analysis will be used to guide the representation of the similar uncertainty in the production plans used in the Sim-Opt framework. Quantifying the stochasticity will make the findings in this study more reliable, as the uncertainty applied in the framework will be as realistic as possible. As described before, weekly data is available. The same holds for the production plans, which were updated every week based on demand forecasts and planning decisions.

4.2.1. Sources of uncertainty

The objective of this analysis was to determine what the source of the observed stochasticity in the production plans is. A company expert suggested that a part of the variability might be caused by planning decisions, such as reallocating production between locations and resources. It is important to know whether this is a valid assumption or if the stochasticity can be fully attributed to demand uncertainty, as this study focuses on applying demand uncertainty in the Sim-Opt framework. Demand variability is implicitly incorporated in production plans, but the plans can have other sources of uncertainty as well. Because of this, the observed variability was decomposed into a demand-driven component and a planning-induced component. For this analysis, only products that are produced on multiple location-resource combinations were taken into account.

$$\mathrm{Var}_{\text{demand}} = \frac{1}{T} \sum_{t=1}^{T} \left(Q_t - \bar{Q}\right)^2 \tag{4.1}$$

The demand-driven variance measures the variability in the total weekly production volume of a product, and is calculated with equation 4.1. Here, Q_t denotes the total quantity planned for production in week t, while $\bar{Q}$ represents the average weekly production volume over the entire observation period consisting of T weeks. This metric captures fluctuations in the overall production level and is interpreted as the variability originating from changes in demand. If the weekly production volume varies substantially over time, the demand-driven variance will be large.

$$\mathrm{Var}_{\text{planning}} = \left(\bar{Q}\right)^2 \cdot \sum_{l=1}^{L} \left[\frac{1}{T} \sum_{t=1}^{T} \left(s_{l,t} - \bar{s}_l\right)^2 \right] \tag{4.2}$$

The planning-induced variance measures the variability caused by changes in the allocation of production across location-resource combinations. Equation 4.2 shows how it is calculated. Let $s_{l,t}$ denote the share of the total weekly production volume assigned to location-resource combination l in week t, and let $\bar{s}_l$ denote its average share over the observation period. The variance of these shares reflects the degree to which production is reallocated between locations and resources from week to week. Multiplying the share variance by the square of the average production volume, $\bar{Q}^2$, converts the variance in dimensionless shares into a variance expressed in production quantities. This makes it comparable to the demand-driven variance.

$$\text{Demand share} = \frac{\mathrm{Var}_{\text{demand}}}{\mathrm{Var}_{\text{demand}} + \mathrm{Var}_{\text{planning}}} \tag{4.3}$$

$$\text{Planning share} = \frac{\mathrm{Var}_{\text{planning}}}{\mathrm{Var}_{\text{demand}} + \mathrm{Var}_{\text{planning}}} \tag{4.4}$$

For each product, a demand share and a planning share are calculated with equations 4.3 and 4.4 respectively. The demand share represents the proportion of total observed variability that can be attributed to fluctuations in total production volume The planning share represents the variability belonging to changes in production allocation.

Table 4.1: Shares of observed variability in production plans

Metric	Value
Average demand-driven variance share	0.9998404
Average planning-induced variance share	0.0001596

The results presented in Table 4.1 show that the average demand-driven variance share equals 0.9998404, while the average planning-induced variance share equals 0.0001596. This implies that more than 99.98% of the observed variability in production plans can be explained by changes in total production volume, leaving less than 0.02% attributable to planning decisions. Consequently, the allocation of production across location-resource combinations remains highly consistent over time and contributes negligibly to the overall stochasticity observed in the production plans. This outcome implies that the production planning process itself is very stable over time. This finding is important because it justifies treating the uncertainty as primarily demand-driven and it eliminates the need to introduce an additional stochastic layer for planning decisions.

4.2.2. Demand uncertainty

As the stochasticity can be attributed almost solely to fluctuations in demand, the next objective is to quantify this uncertainty. From here onwards, when referring to uncertainty or something similar, we refer to demand uncertainty, unless stated otherwise. Due to the restrictions in the data availability, the stochasticity is analysed over a period of 6 months. This means that certain trends in demand, caused by holidays or weather for example, are not fully represented in the analysed production plans. It is however, important to get an idea of what the stochasticity looks like across a planning horizon of a full year. Therefore, a method based on block bootstrapping was used to analyse the data. Block bootstrapping is a resampling technique used in statistics to analyse serially correlated data [40]. The data is cut into contiguous blocks, yielding independent samples. This method preserves the internal correlation structure within each block and can identify regime behaviour. Figure 4.1 shows an example of how 10 consecutive production plans would be analysed using this technique. By using this method, relative blocks are used to analyse demand stochasticity, instead of absolute blocks which is impossible due to the data restrictions.

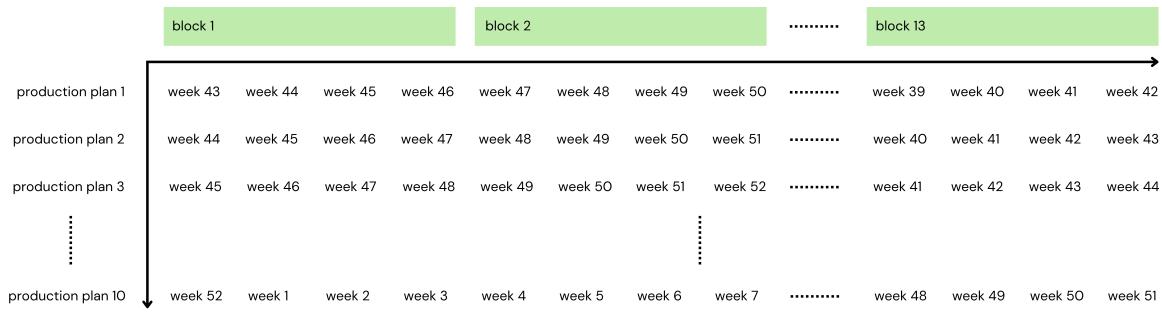

Figure 4.1: Production plan block bootstrapping

For every location-resource-product combination in the production plans, the production volume across each 4-week blocks is calculated. The range of 4 weeks per block was chosen to account for local changes, such as shifting volume to the neighbouring week. For each block, the deltas between consecutive production plans are determined. The goal is to identify how large these deltas are and how they are distributed, to understand how much the production volumes change across the planning horizon. With this information, a probability density function (PDF) is fitted for every block, which describes the distribution of the deltas in the data. One calculation of a delta between two production plans for block 1 looks as follows:

$$\Delta v^1 = v^1_{\text{current}} - v^1_{\text{previous}} \tag{4.5}$$

Where:

- Δv^1 = change in production volume
- v^1_{current} = production volume of block 1 in the current production plan
- v^1_{previous} = production volume of block 1 in the previous production plan

To create as many data points as possible to fit each PDF, the analysis is performed on location-resource-product level. This means that the volume per block is calculated for every location-resource-product combination. The deltas are therefore also determined at this level, providing as much data as possible. To determine a representable delta for each block, the average of all deltas between consecutive production plans in that block was calculated. This, together with the standard deviation of the deltas, is shown in Figure 4.2 for each block.

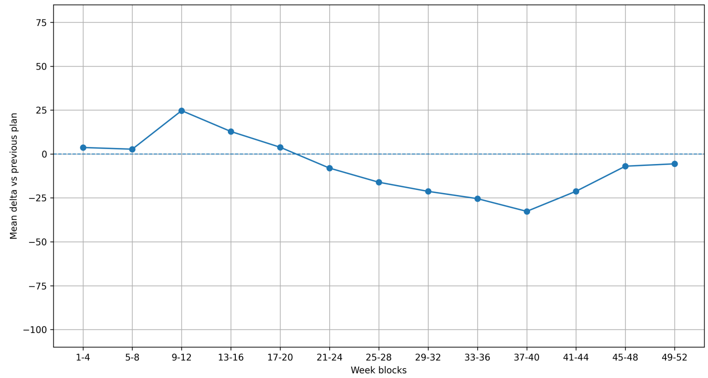

(a) Mean of deltas per rolling 4-week block

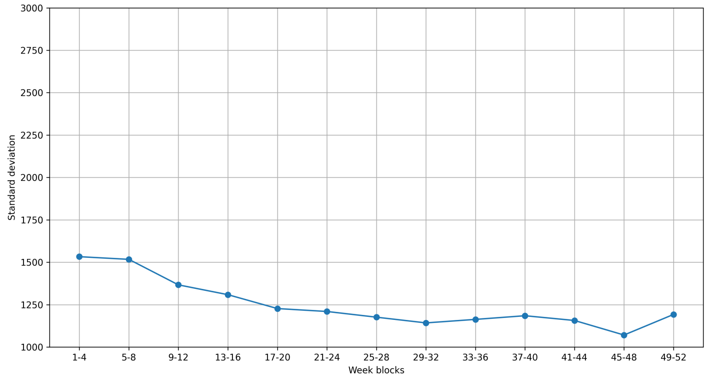

(b) Standard deviation of deltas per rolling 4-week block

Figure 4.2: Metrics of **all** deltas per rolling 4-week block

These results give an idea about the expected delta between consecutive production plans, together with the variation of the deltas in each block. The higher standard deviations in the first few blocks suggest that a lot of demand uncertainty is present in the shorter planning horizon. This could be explained by changes in production week between the blocks. The positive deltas in the first blocks, in contrast to the negative deltas in later blocks, also indicate that production levels are increased in the shorter horizon, while decreased on a medium to long planning horizon. Meanwhile, these results might also indicate a seasonal pattern that overfitted due to the data restrictions. To fit the data of each block to a PDF, kernel density estimation (KDE) is used. KDE is a non-parametric method used to estimate the underlying probability density function of a random variable based on a finite set of data points [41]. Instead of assuming a specific distributional form (like a normal distribution), KDE builds the density estimate directly from the data. It does this by placing a smooth, symmetric function called a kernel (e.g. a Gaussian) centred at each data point. The estimated density at any location is obtained by summing the contributions of all these kernels. This results in a smooth curve that reflects where data points are concentrated. Using KDE was necessary to provide a good distribution for this data.

Table 4.2: Shares of observed (non)zero deltas in **block 1**

	Observations	Share
Entire set of deltas	118896	1
Zero values	85844	0.72
Nonzero values	33052	0.28

Due to the shear number of blocks for which a PDF will be fitted, block 1 will be used as an example. Figure 4.3 shows a histogram of the delta distribution in block 1, together with the PDF fitted by the KDE method. The distribution of the deltas goes to several extremes. The histogram shows that there is a huge concentration of deltas around zero. The PDF shows a very tall and narrow spike near zero as well, representing the high density at the centre. The share of observations in which delta is zero is very high. Table 4.2 shows the distribution of (non)zero deltas which were observed. It states that 72% of all delta observations in this block is zero, which substantiates what can be seen in Figure 4.3. This indicated that the deltas can be described as zero-inflated data. What is barely visible in the histogram, is that there are several outliers up until -90.000 and +90.000, which gives the PDF its long and thin tails. Due to these extreme outliers, and the high number of zero observations, the central region looks very compressed and is not very visible.

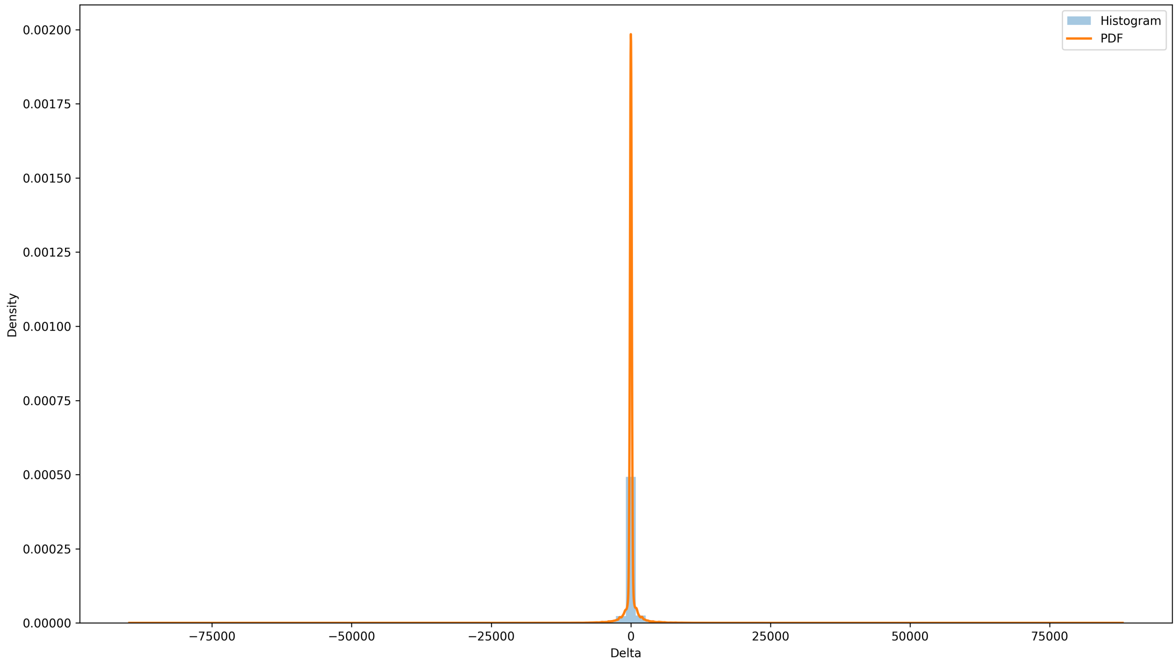

Figure 4.3: Delta distribution with PDF of **block 1**

Each 4-week block over a planning horizon of one full year has a unique distribution of (non)zero deltas in that block. Table 4.3 shows the distributions of each block.

Table 4.3: Shares of observed (non)zero deltas in all blocks B1 to B13

Delta	B1	B2	B3	B4	B5	B6	B7	B8	B9	B10	B11	B12	B13
Zero	0.72	0.68	0.67	0.64	0.64	0.65	0.65	0.66	0.67	0.67	0.68	0.69	0.70
Nonzero	0.28	0.32	0.33	0.36	0.36	0.35	0.35	0.34	0.33	0.33	0.32	0.31	0.30

4.2.3. Zero-truncated delta distribution

All blocks have similar shares of (non)zero delta observations. Due to the high share of zero observations in all blocks, it was decided to continue fitting a probability density function to only nonzero values. This way, it would be possible to get a better idea of the actual distribution of deltas and this would make it more visible as well. This split could be described as a hurdle formulation where the nonzero distribution proposed in this section is zero-truncated [42]. The same analysis was performed as in Section 4.2.2, but now simply only taking into account nonzero values. Figure 4.4 shows the observed deltas, together with the standard deviation of these deltas, for each block. While the graphs show similar trends, compared with Figure 4.2, the average deltas are much higher and their standard deviations as well. This is logical due to the fact that all zero values are taken out of the equation.

Figure 4.5 shows a histogram of the zero-truncated delta distribution in block 1, together with the PDF fitted by the KDE method. The plot shows that there is still a high concentration of deltas around the centre, but much less than in Figure 4.3. The PDF also still has a very high density near the centre, but the peak did became less narrow. Its tails are still very flat due to several extreme outliers. The analysis shows that most of the observations are between -10.000 and +10.000. Appendix D shows the zero-truncated distributions with PDFs for all 13 blocks across the horizon.

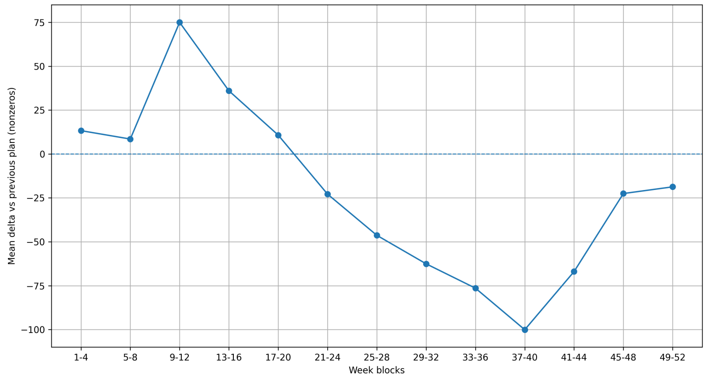

(a) Mean of **nonzero** deltas per rolling 4-week block

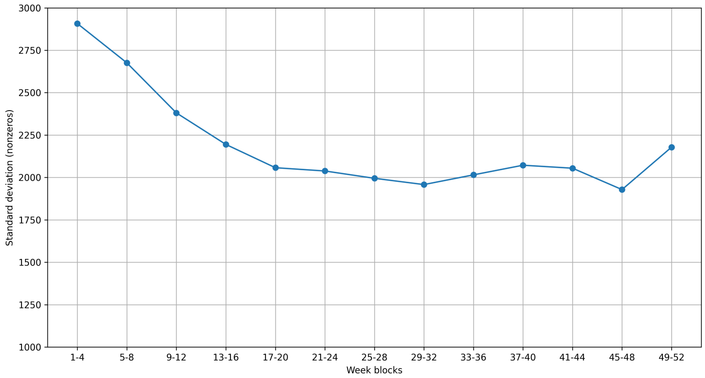

(b) Standard deviation of **nonzero** deltas per rolling 4-week block

Figure 4.4: Metrics of **nonzero** deltas per rolling 4-week block

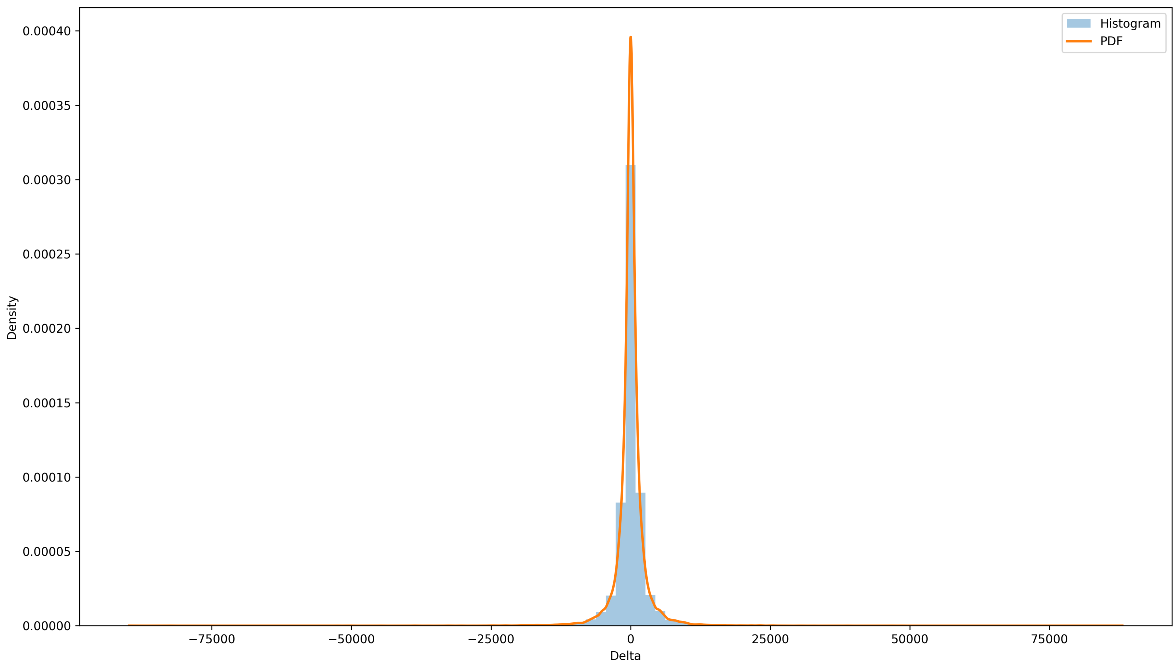

Figure 4.5: **Zero-truncated** delta distribution with PDF of **block 1**

4.2.4. Subcategories

To see if a differentiation in the stochasticity of certain products exists, the same stochasticity analysis is performed on some product groups. Table 4.4 shows the different product groups by CMG subcategory code, and how many products are observed in all production plans and with that code. The same analysis was performed as in Section 4.2.2, to see what the deltas between consecutive production plans is across the planning horizon.

Table 4.4: Occurrence of product types in all production plans (U-Z represent product types)

CMG subcategory	Product group description	Observations
GA02	Finished products - U	3884
FA01	Semi-finished products - U	665
GA03	Finished products - V	298
FA02	Semi-finished products - V	38
GA11	Finished products - W	28
GA07	Finished products - X	24
GA06	Finished products - Y	14
GA10	Finished products - Z	3
Total		4954

The five smallest subcategories represent only 2.2% of the total number of products, and are therefore not taken into account. Their analysis results can be found in Appendix E. As the remaining product groups GA02, FA01 and GA03 represent the vast majority of all products, these categories are important to compare. Figure 4.6 shows the analysis results of the GA02 and FA01 subcategories. Subcategory GA02 shows a similar pattern as the overall results in Figure 4.4. While the behaviour of deltas for subcategories FA01 and GA03 (Appendix E) does follow a different pattern, this is mostly caused by outliers. An additional breakdown of the production plans into these subcategories is not expected to lead to major improvements, as subcategory GA02 is by far the largest group. Though slight improvements in the accuracy of the PDFs could possibly be found, it is unlikely that the representation of the true stochasticity in the production plans would improve significantly. Therefore, it is decided to continue this research with the distributions and PDFs from Section 4.2.3.

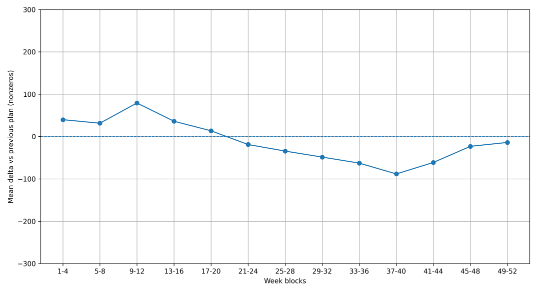

(a) Mean **(GA02)**

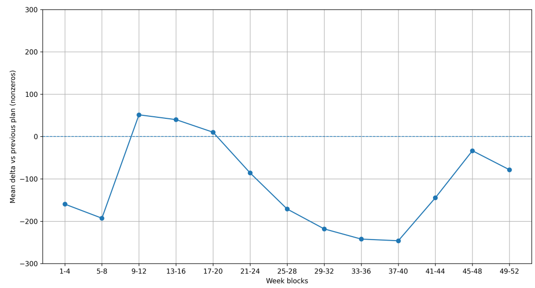

(b) Mean **(FA01)**

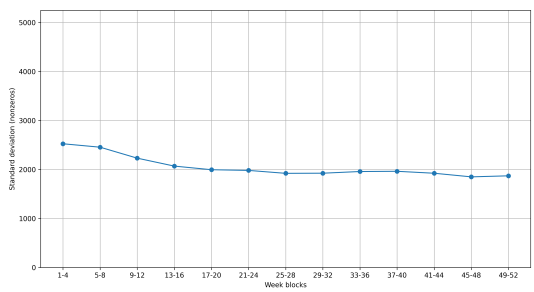

(c) Standard deviation **(GA02)**

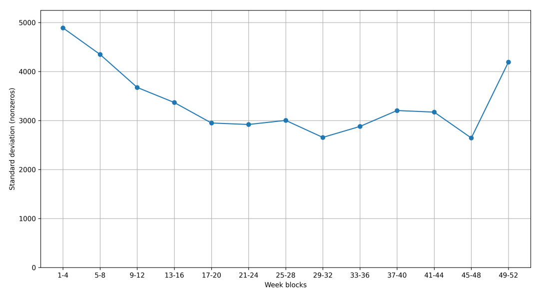

(d) Standard deviation **(FA01)**

Figure 4.6: Metrics of **nonzero** deltas per rolling 4-week block for subcategories GA02 & FA01

4.3. Material analysis

To find changeover materials for which safety stock can be especially beneficial, a material analysis was performed. This analysis is performed on an item-level, so each individual SKU is analysed. The goal of this analysis is to be able to classify items based on their characteristics. This analysis focuses on the material demand and the value they represent in the manufacturer's inventory management.

4.3.1. Categorisation of demand patterns

The changeover materials can be classified according to different demand patterns. Syntetos et al. proposed a categorisation scheme based on the average inter-demand interval (ADI) and the square of the coefficient of variation (CV^2) [43]. Their study aims to categorise demand patterns to facilitate the selection of forecasting methods. While using the classification for a different purpose in this study, it is still very useful. Figure 4.7 shows how material demand can be divided into four patterns; erratic, lumpy, smooth and intermittent. The figure also shows what the demand behaviour of these patterns look like. On the x-axis, the ADI describes the variability in demand timing. Meanwhile, the CV^2 describes the variability in demand quantity on the y-axis. According to Syntetos et al., the cut-off values at which the demand behaviour changes to another pattern are 1.32 for ADI and 0.49 for CV^2 [43].

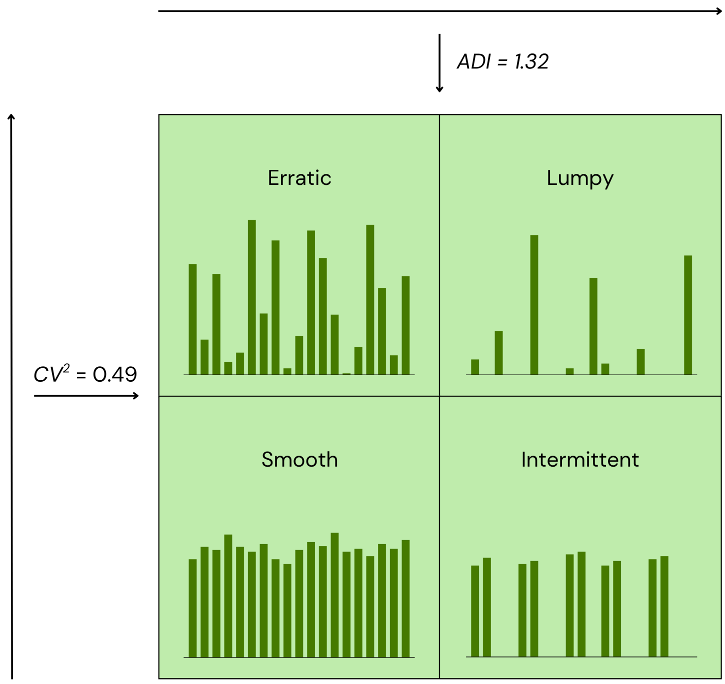

Figure 4.7: Demand pattern classification method

To perform the analysis on the changeover materials used in the manufacturer's supply chain, the ADI and CV^2 are calculated according to equations 4.6 and 4.7 respectively.

$$ADI = \frac{\text{\# periods}}{\text{\# demand buckets}} \tag{4.6}$$

where # periods is the total number of weeks and # demand buckets the number of weeks in which there is a demand.

$$CV^2 = (\frac{s}{\bar{X}})^2 \tag{4.7}$$

where s is the sample standard deviation and $\bar{X}$ the sample mean.

Figure 4.8 shows an enlarged view of the demand pattern classification results. The red lines represent the classification boundaries defined by Syntetos et al., with an ADI threshold of 1.32 and a CV^2 threshold of 0.49 [43]. Each point represents an item in the manufacturer's supply chain. The results show that there are 35 items with erratic demand, 941 with lumpy demand, 247 with smooth demand, and 4607 with intermittent demand behaviour.

The distribution of items across the four demand categories reveals that the inventory is dominated by items with irregular demand characteristics. More than 79% of all items are classified as intermittent. This shows that demand occurs infrequently and is separated by relatively long periods without consumption. However, when demand does occur, the quantity remains relatively stable, as shown by the low CV^2 values. This suggests that the primary source of uncertainty for most changeover materials is the timing of demand rather than the quantity required.

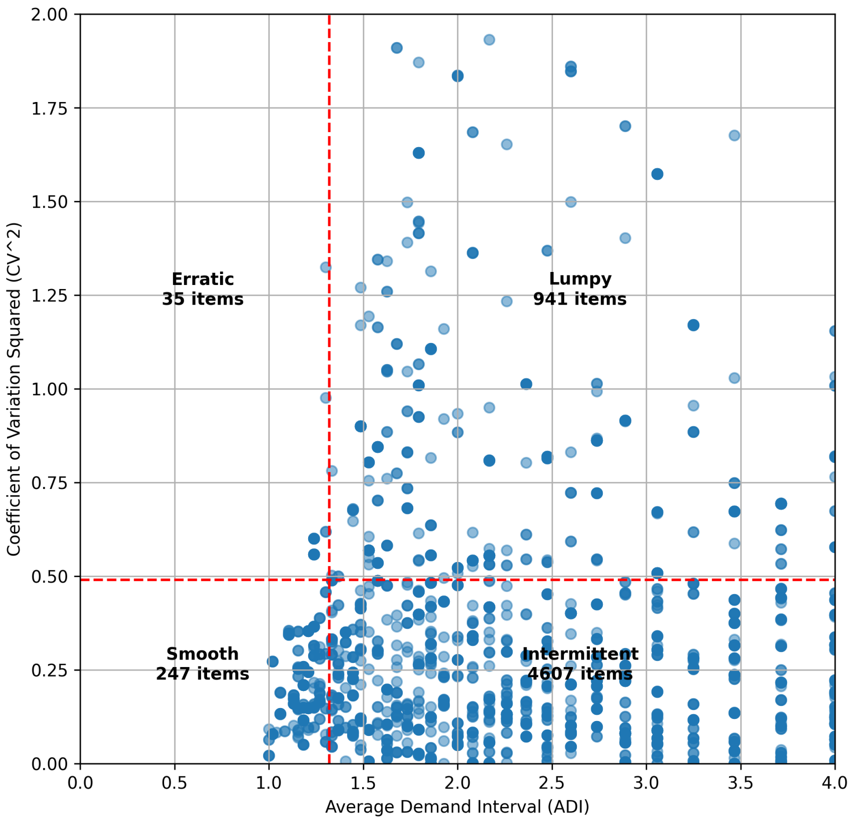

Figure 4.8: Material demand classification of changeover materials (enlarged)

A further 16% of the items are classified as lumpy. These items exhibit both infrequent demand and large variability in demand quantities. Lumpy items are generally the most difficult to plan and forecast, as uncertainty exists in both timing and quantity. A substantial portion of the inventory is therefore exposed to high forecasting uncertainty. Only a small proportion of items are classified as smooth (4%). These items are characterised by frequent demand and low variability in demand quantities, making their consumption behaviour relatively predictable. Similarly, only 35 items are classified as erratic. Although these materials are demanded frequently, the quantities fluctuate considerably between demand periods. The low shares of smooth and erratic items suggest that regularly consumed materials account for a minor part of the changeover materials.

The scatter plot in figure 4.8 illustrates a strong concentration of items for which ADI values exceed the threshold of 1.32. This indicates that demand intermittency is a dominant characteristic of the manufacturer's changeover materials. As a result, inventory policies and forecasting approaches should primarily focus on managing intermittent and lumpy demand behaviour.

4.3.2. ABC analysis

The ABC inventory classification model is used to classify items according to their importance. ABC analysis is based on the Pareto principle, also known as the 80/20 rule, which states that a relatively small proportion of items is responsible for the majority of the total value generated by the inventory [44]. In inventory management, this principle implies that only a limited number of stock keeping units (SKUs) contribute significantly to revenue. Meanwhile, a large number of items contribute relatively little. As a result, inventory items are divided into three classes according to their importance: class A contains the most critical items, class B contains items of moderate importance and class C contains the least critical items and therefore require the lowest level of management attention [44]. According to the literature, ABC classifications typically follow a Pareto distribution in which class A represents approximately 10–20% of the items, class B around 20–30%, and class C the remaining 50–70% of the inventory [44].

Figure 4.9 presents the Pareto distribution obtained from the material demand data used in this study. The cumulative value curve demonstrates the highly uneven contribution of items to the total inventory value. The results show that only 13.9% of all items (813 items) account for 80% of the total value and are therefore classified as class A. These items represent the most economically important part of the inventory and should receive the highest level of control and replenishment attention.

The next 20.1% of items (1 168 items) increase the cumulative contribution from 80% to 95% of the total value and are assigned to class B. Although these items are less critical than class A, they still contribute substantially to overall inventory performance and require regular monitoring. The remaining 66.0% of material demand (3,849 items) contribute only 5% of the total value and are classified as class C. Despite representing the majority of stock items, their individual impact on the total inventory value is relatively low. These results closely follow the Pareto principle and confirm a strong concentration of inventory value within a small subset of items.

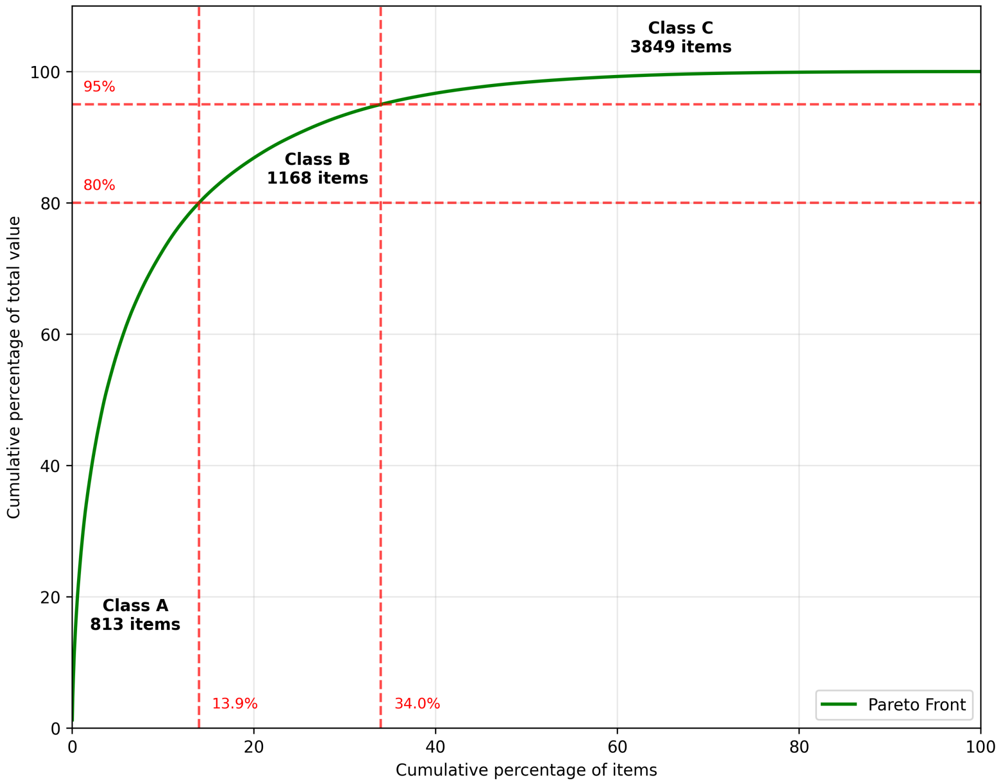

Figure 4.9: Pareto distribution of ABC analysis

5 Modelling framework

To understand how the model aims to find results for this research, several design choices have to be explained. While the general idea of the Sim-Opt framework was described in Section 1.6, this chapter will specify in detail how the model is built and used. The model design was based on the performed literature review, current state analysis and data analysis, and will therefore be elaborated further in this chapter. While changes to the model can be made in the experiments, this chapter describes the base design of the model.

5.1. Hard safety stock constraint

As mentioned in Section 3.2.1, the safety stock constraint currently used in the MRP-Opt model is considered a soft constraint. This means that a slack variable is added to the constraint, which allows the material stock to fall below the predetermined safety stock level. To parametrise safety stock levels of changeover materials, it is crucial to get rid of this slack variable. Therefore, this research proposes a new parameter, called the hard safety stock (HSS). This is also the parameter which the Sim-Opt framework will optimise. To do this, another safety stock constraint is also added to the MRP-Opt model, shown in Formula 5.1.

- **Hard safety stock constraint:** If a production facility uses a given material in a given week, then the available stock $s_{\ell,m,w}$ at location l, for material m in week w must be at least the required hard safety stock level $HSS_{\ell,m,w}$ of this material. This constraint can be seen as a hard constraint, as the model does not allow stock levels to fall below the established hard safety stock level through a slack variable.

$$s_{\ell,m,w} \geq HSS_{\ell,m,w} \tag{5.1}$$

Although adding a constraint could be seen as a major change to the existing MRP-Opt, it was necessary due to the structure of the objectives. The prioritising nature of the objectives, as described in Section 3.2.1, pushes the optimisation towards minimum obsolete stock costs. The sub-objective minimising obsolete stock costs has a higher priority than the sub-objective minimising the safety stock slack penalty. Having safety stock often means having higher obsolete stock costs, which pale into insignificance with the financial gains created by safety stock in the long run. As these improvements are not taken into account in the MRP-Opt, the model will rarely choose to minimise the safety stock slack penalty before it minimises the obsolete stock costs. With a slack variable present in the safety stock constraint, optimising the safety stock levels through the existing soft safety stock parameter was not possible. When talking about safety stock in the remainder of this thesis, we refer to the hard safety stock parameter.

5.2. Optimisation parameters

As mentioned, the goal of this research is to optimise safety stock levels. To find these safety stock levels as accurately as possible, it is useful to express the safety stock in terms of **days of sales** (DOS). One-day-of-sales is defined as the number of materials in stock which is needed to satisfy 1 day of demand, or sales. Therefore, the safety stock levels can be expressed in days of sales. Expressing the safety stock in days of sales provides more precise safety stock determination, as all materials could have a safety stock of e.g. 10 days of sales, but the absolute safety stock value for each material is different. To create the safety stock values this way, the average daily demand (1 DOS) was calculated for each material. This creates buffers that are more realistic compared to setting an absolute safety stock value to all materials. Due to the sheer number of materials in the manufacturer's supply chain, classifying materials into material groups can provides more feasibility. While the materials are aggregated this way, losing some heterogeneity, it is easier to find HSS values for them. For all materials in a material group, the same HSS parameter in terms of days of sales will be optimised. The classification of these material groups will be explored in Section 5.5. This approach does implicitly assume that materials have a steady material consumption, since 1 DOS is calculated as average daily demand. While in reality, a material's demand pattern can be very lumpy or intermittent, as shown in Section 4.3.1. The demand for some materials is not steady, making the 1 DOS value to be less meaningful than for materials with steady demand. Other metrics, such as safety time, might be more useful to quantify the safety stock for these materials. While noting this limitation, the days of sales approach will be tested in this research.

To find the optimal values for the safety stocks within reasonable computation time, sampling bounds are assigned to each material group. Sampling bounds make sure the model does not test any extreme solutions. The bounds are the same for each material group, as all materials have a unique average daily demand. The lower bound is zero, as it should always be possible to operate without safety stock. The upper bound is set to 7 days of sales. This is a reasonable starting point, as practice-oriented literature often suggests maintaining buffer levels on the order of a few weeks of demand (e.g. two weeks or more, depending on uncertainty) [45]. Depending on the first results, the bounds could be expanded to more days of sales.

5.2.1. Hard Safety Stock assignment variants

The hard safety stock will be assigned to the MRP-Opt tool on a material-location, or item, level. When the Sim-Opt framework is run for a test set of items, the material-location combinations will be extracted from this test set. The HSS will then only be applied to these material-location combinations. The hard safety stock can be assigned by the Sim-Opt framework to the MRP-Opt model in three different ways:

- **Constant variant:** In the constant HSS formulation, a hard safety stock policy becomes active from the simulation start week and remains active for all subsequent weeks in the MRP-Opt horizon, which is 52 weeks. Methodologically, this represents a stationary buffering regime, allowing to isolate the effect of a constant inventory protection level over the full horizon. Because activation is permanent after initialization, this variant is well suited as a baseline scenario for comparison against more dynamic policies.
- **Changeover weeks variant:** In the changeover weeks HSS formulation, each material-location pair receives an item-specific activation window bounded by the minimum and maximum changeover week. The HSS is only applied to the weeks within that closed interval. This corresponds to a finite-horizon protection strategy, allowing heterogeneous activation timing across items. This improves buffer accuracy when demand uncertainty is concentrated in specific periods rather than uniformly distributed across the horizon.
- **Max changeover week variant:** In the maximum changeover week HSS formulation, the HSS is only activated in the final changeover week, creating a one-week intervention for each item specifically. Conceptually, this is a buffering policy designed to protect against a narrowly localised risk moment, such as an anticipated disruption or demand spike tied to a specific calendar week.

5.3. Objective function and decision variable

The objective function on which the Sim-Opt will learn and optimise the safety stock levels will be based on the obsolete stock decision variable. Obsolete stock is left over material stock due to a material changeover. As mentioned in Section 3.2.1, the MRP-Opt model provides the obsolete stock values for each material as output of the MIP. The decision variable is multiplied by the effective material cost, creating the optimisation objective function in Formula 5.2.

- **Optimisation objective function:** The obsolete stock $obs_{\ell,m,w}$ of material m at location ℓ in week w is multiplied by the piece cost c_m of material m. The objective function minimises the sum obsolete stock costs over all materials and locations in week w.

$$\min \sum_{\ell \in L} \sum_{m \in M} c_m \cdot obs_{\ell,m,w} \tag{5.2}$$

The simulation-based optimisation will use this objective function to determine the safety stock levels for the different material groups. Since a simulation run aims to represent reality, it runs multiple sequential weekly MRP-Opt runs which represents the MRP process at the manufacturer. The obsolete stock costs in the final week therefore represent the costs which the manufacturer is stuck with. While each MRP-Opt run generates a unique objective value, the objective of the final week in the sequential process is therefore the value used in the safety stock optimisation. The Sim-Opt framework was illustrated in Section 1.6. Due to the stochastic nature of the model, multiple simulation runs are performed to account for the demand variability. Because of this, the average objective value of all simulation runs in a Sim-Opt iteration is calculated and used to learn the safety stock parameters. Figure 5.1 shows an example of one Sim-Opt iteration, with multiple simulations and 4 MRP-Opt runs in each simulation. The objective function value of the final MRP-Opt run is picked as the objective value for a simulation, and the average of all these simulation objectives represents the weighted objective value of the optimisation iteration.

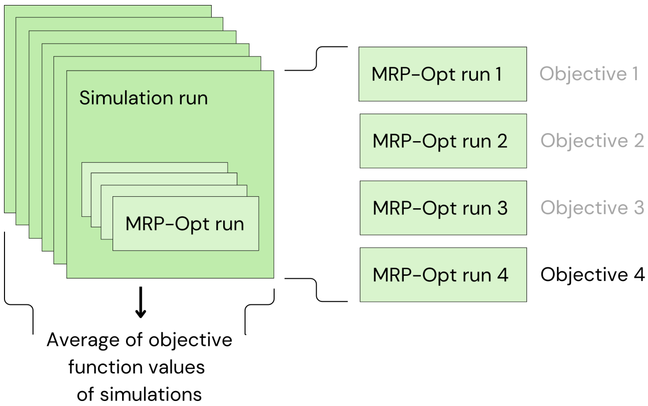

Figure 5.1: Objective function value structure example

As mentioned in Section 1.1, adding safety stock to a MRP environment creates a trade-off between safety stock (inventory) costs, obsolete material costs and rush order costs. To parametrise safety stock levels as realistically as possible, this trade-off should preferably be represented in the objective function. Adding inventory costs and rush order costs to the objective function would make it a multi-objective function. Formula 5.3 shows these additions to the objective function, where the inventory cost $v_{\ell,m}$ is multiplied by the stock level $s_{\ell,m,w}$ and a rush order penalty $p_{\ell,m}$ is multiplied by the necessary material order $in_{\ell,m,w}$.

$$\min \sum_{\ell \in L} \sum_{m \in M} (c_{\ell,m} \cdot obs_{\ell,m,w}) + (v_{\ell,m} \cdot s_{\ell,m,w}) + (p_{\ell,m} \cdot in_{\ell,m,w}) \tag{5.3}$$

Due to a lack of data generated by the MRP-Opt model, the use of this multi-objective function is not feasible for this research. Data on inventory costs and costs regarding rush orders, is not generated by the model. Therefore it was decided to choose the above mentioned objective as a means of representing this trade-off as best as possible. Another expansion of the objective function was considered. This expansion has to do with penalising material shortages. Currently, the objective function penalises being left with obsolete materials. As setting safety stock levels tries to prevent material shortages when material demand spikes due to uncertainty, it could be interesting to penalise material shortages in the objective function as well. However, due to the nature of the MRP-Opt model, which does not allow material shortages, this is not a possible expansion. The MIP is built to never allow material shortages at production facilities, which would be necessary to cater this addition. However relevant for literature, it is not feasible for this research.

5.4. Optimisation method

A **Full Enumeration** method is proposed as the approach for exploring the safety stock parameter space. This method evaluates every possible combination of parameter values within the predefined search domain. As discussed in Section 5.2, the this search domain is bounded by 0 and 7 days of sales. This search domain will be explored for every integer value. By systematically testing all candidate solutions, it guarantees that the globally optimal solution will be identified, provided that the search space is finite and fully accessible. However, the computational cost grows rapidly with the number of optimisation parameters. Full Enumeration is therefore mostly practical for problems with a relatively small number of parameters or limited parameter ranges. Therefore, this method is chosen to explore the HSS parameter space as defined in Section 5.4. Because no regions of the parameter space are ignored, Full Enumeration can serve as a reference against which other optimisation methods are compared in future research.

5.5. Material group classification

To find materials for which safety stock can be beneficial, a material classification was performed. This classification is based on the analyses performed in Section 4.3. The goal is to classify materials into material groups (MGs) and to identify which MGs are most promising to test in the Sim-Opt framework. The combination of the ABC classification and the demand pattern classification results in twelve distinct material groups, as shown in Table 5.1. This two-dimensional classification clusters materials according to their economic importance and demand characteristics. Combining both dimensions allows for a more targeted evaluation of safety stock policies.

Table 5.1: Material group classification

Material group [–]	**ABC class** [–]	**Demand pattern** [–]	**Group size** [items]
MG1	A	Erratic	19
MG2	A	Lumpy	100
MG3	A	Smooth	78
MG4	A	Intermittent	457
MG5	B	Erratic	17
MG6	B	Lumpy	153
MG7	B	Smooth	77
MG8	B	Intermittent	748
MG9	C	Erratic	12
MG10	C	Lumpy	480
MG11	C	Smooth	72
MG12	C	Intermittent	2 687

A clear observation from Table 5.1 is that the majority of materials are concentrated in the intermittent demand categories (MG4, MG8 and MG12). This is consistent with the results presented in Section 4.3.1, where intermittent demand was identified as the dominant demand pattern within the manufacturer's portfolio. In particular, MG12 contains 3 171 items, making it the largest material group by far. Although these materials belong to class C and therefore contribute relatively little to the total inventory value, their large population makes them relevant when evaluating inventory management policies. From a safety stock perspective, however, the most promising material groups are not necessarily the largest groups. Safety stock creates the greatest benefit for materials that simultaneously have a significant business impact and exhibit demand uncertainty. Because of this, the material groups belonging to class A are of particular interest because these materials account for the largest share of the total inventory value. These materials therefore have the greatest potential impact on inventory costs.

Within class A, MG2 (A-Lumpy) and MG4 (A-Intermittent) appear especially promising for further analysis. Materials in MG2 are characterised by both infrequent demand and high variability in demand quantities. This combination makes demand difficult to forecast and increases the risk of stockouts. As a result, safety stock levels may significantly improve material availability while reducing disruptions caused by unexpected demand peaks. MG4 contains materials with intermittent demand, where demand timing is uncertain but quantities are relatively stable when demand occurs.

MG3 (A-Smooth) is expected to be less suitable for extensive safety stock optimisation. The relatively predictable demand behaviour of smooth materials reduces forecasting uncertainty. Replenishment planning can generally be performed accurately without requiring substantial safety stock. Similarly, MG1 (A-Erratic) contains only a limited number of materials, making it less representative for generalisable conclusions regarding safety stock policies. The B-class material groups form an intermediate category. While these materials are less critical than class A materials, groups such as MG6 (B-Lumpy) and MG8 (B-Intermittent) may still benefit from safety stock due to their demand uncertainty. The C-class material groups are generally expected to offer the lowest return from safety stock optimisation. Although groups like MG10 (C-Lumpy) and MG12 (C-Intermittent) contain a large number of materials, their limited value contribution reduces the financial impact of stockouts.

5.6. Items in the Sim-Opt framework

A material can be used at multiple locations, or used for multiple products at the same location. Therefore, the data should be processed with that distinction. These material-location-product combinations are called **items**. These items could also be described as Stock Keeping Units (SKUs). Much of the data is available on an item level. The HSS parameters will therefore also be applied to items, not to materials. While items and materials can sometimes by used simultaneously, there is a clear distinction. Table 5.2 shows an example dataset of relevant data for 5 randomly chosen items. This dataset consists of gathered data from the manufacturer's MRP-Opt tool, together with data generated in this research. The data is available per item, meaning for every unique combination of material-location-product. When running the Sim-Opt framework, items are filters on a specific material group, so the framework will only test with items which have that material group.

Table 5.2: Example item dataset (empty rows contain classified information)

Parameter	Unit	Item/SKU				
		1	2	3	4	5
Material	[–]					
Location	[–]					
Product	[–]					
Vendor	[–]					
MOQ	[units]	96 200	80 028	80 028	114 400	114 400
Piece cost	[€/unit]	0.1415566	0.11208	0.11208	0.084	0.084
CMG subcategory	[–]	GA02	GA02	GA02	GA02	GA02
Total demand	[units/year]	3 485 651	453 165	745 048	13 576 376	4 438 601
1 day of sales	[units/day]	9 550	1 242	2 041	37 196	12 161
Changeover	[–]					
Project	[–]					
Min week	[YYYY-WW]	2025-47	2026-10	2026-10	2025-49	2025-49
Max week	[YYYY-WW]	2025-47	2026-20	2026-20	2026-09	2026-09
Expected week	[YYYY-WW]	2025-45	–	–	–	–
Material group	[–]	MG4	MG8	MG6	MG4	MG2
Demand pattern	[–]	Intermittent	Intermittent	Lumpy	Intermittent	Lumpy
ABC class	[–]	A	B	B	A	A

5.7. Uncertainty assignment

Uncertainty will be added to the simulation by manipulating the production plan in a sequential way. By running the manufacturer's MRP optimisation model week by week, while manipulating the production plan every week, a stochastic production environment will be imitated. Section 4.2.2 describes the stochasticity in the manufacturer's production based on historical data. To add **realistic** stochasticity to the production plans in the Sim-Opt procedure, the variabilities found in the data analysis in Chapter 4 will be applied. Just like in the uncertainty analysis, the assignment of uncertainties will be done based on the block bootstrapping method. To create a fair comparison between MRP-Opt results, each MRP-Opt run will optimise the changeovers over a horizon of 52 weeks. This means that the production plan given to the MRP-Opt model will have 52 weeks. Therefore, the uncertainty will also be assigned over a 52 week period, with 13 blocks of 4 weeks. The assignment of variability to the production plans will be explained following the 4 levels revealed in Figure 5.2.

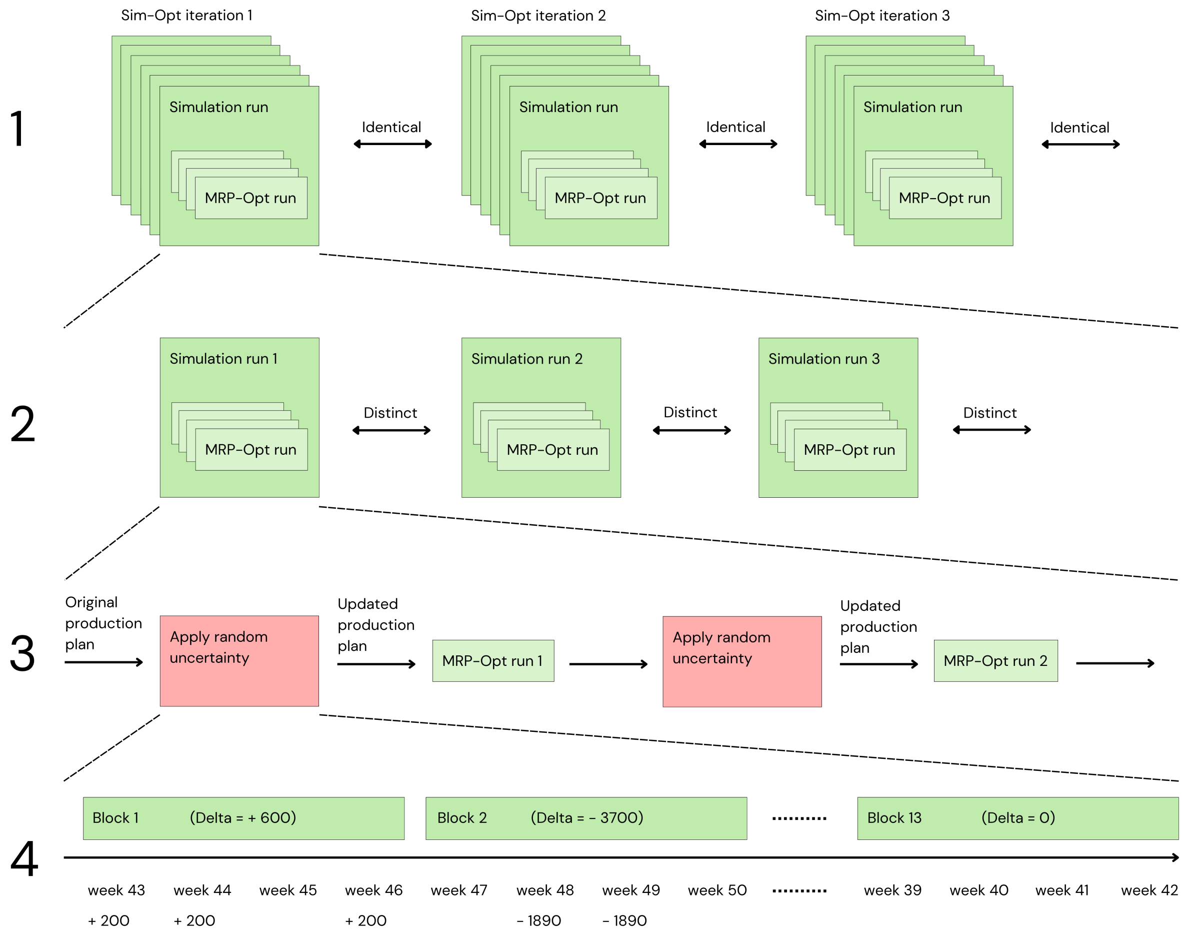

Figure 5.2: Realistic uncertainty assignment levels in the Sim-Opt procedure

At the start of the Sim-Opt procedure, a historical production plan will be picked. This production plan contains real data from a certain week in which it was active at the manufacturer. This week is therefore also the starting week, meaning the first MRP-Opt run will run for that week in the MRP-Opt model. Uncertainty will be applied to the production plan **before** each MRP-Opt run in a simulation. **Level 3** in figure 5.2 shows this process. Random uncertainty is applied to the production plan, after which the updated production plan is used for the next MRP-Opt run. This production plan will have shifted one week ahead, and will contain new production volumes based on the assigned uncertainty.

The uncertainty will be assigned to the production plan on a location-resource-product level. Meaning the production plan contains production values for each product on a resource at a location. As the found uncertainty in Section 4.2.3 is expressed in terms of a delta Δ between production plans, this is also how the uncertainty will be applied to these plans again. The deltas are sampled from zero-truncated probability density functions (PDFs) from the data analysis. These PDFs are different for each block, therefore the sampled deltas are also different for each block. For each block the following process is executed.

1. A hurdle model is used to apply the deltas to the production plans. A hurdle model consists of a zero-truncated distribution component, such as the PDFs used in this research. The other component is the hurdle component which models the probability of having a (non)zero value [42]. The probability of assigning a nonzero delta is stated in 4.2.2 for each block. This probability determines if a delta will be assigned to this block. If it is decided that a delta will be applied, the value of the delta will be sampled according to the PDF of this particular block. Formula 5.4 shows how this step looks like for block 1. The probability of applying a delta is 0.28. When a delta will be applied, a value is sampled from block 1's PDF.

$$\Delta = \begin{cases} 0, & \text{with probability } 0.72 \\ x, & \text{with probability } 0.28 \end{cases} \tag{5.4}$$

where the random variable x has probability density function $f(x)$ for this block.

2. Looking at Figure 5.2, **level 4** shows that a delta of +600 is sampled for block 1. Now this value has to be distributed among the 4 weeks in this block. The Sim-Opt framework randomly determines among how many of the 4 weeks this delta is distributed, and which weeks this will be. In the case of block 1 in Figure 5.2, the delta is distributed among week 43, 44 and 45. These values will then be added to all production values in that column of the production plan.

This process is repeated for each block, until the full production plan is updated.

Level 2 in Figure 5.2 shows that for each simulation run within the **first** Sim-Opt iteration, different production plans are used. The process of applying uncertainty in the form of deltas to production plans as described by levels 3 and 4, is performed in each simulation of the first iteration. Running multiple simulations with different stochasticity is a core function of the Sim-Opt framework. This way, multiple imitations of reality can be tested.

In contrast to level 2, **level 1** indicates that identical production plans are used for each Sim-Opt iteration. This is to let the safety stock optimisation test its safety stock parameters with the exact same conditions in each iteration. All production plans that are randomly adapted throughout the first Sim-Opt iteration, are saved and used in the following Sim-Opt iterations.

5.7.1. Artificial uncertainty variants

To test the parametrisation of HSS in a more artificial setting, the assignment of uncertainty in the form of deterministic perturbations is explored. A setting in which HSS can be most helpful, is when an unexpected increase in production quantity happens. When no HSS is present as a buffer, high MOQs can result in obsolete materials when the increase in production is not as large as the order quantity. In this case, HSS can be very useful. For changeover materials, it is therefore most relevant to create this scenario around the changeover, as the MRP-Opt tries to end up with zero obsoletes after the changeover. To create this scenario, a sudden positive perturbation is added to the production quantity of the next simulation week. So, before the MRP-Opt model is run,the artificial perturbation is added to the production plan. The production plans are available on a product-location-resource level. When the Sim-Opt framework is run for a set of items, the product-location combinations will be extracted from this test set. Then, the perturbation will only be applied to these product-location combinations in the production plan. There are two variants of when the perturbation is applied:

- **Changeover weeks variant:** In the changeover weeks uncertainty variant, an artificial deterministic perturbation with fixed magnitude is applied to the production plan in the next simulation week. For each material, the perturbation is only applied when the next simulation week is one of the items changeover weeks. This creates an artificial bias, where the production value will be higher than expected. This produces a controlled intervention setting in which the timing and size of the disturbance are known. Figure 5.3 gives an example of how this perturbation is applied to the production plan. For a given material, the changeover must happen between week 44 and 46. When the next simulation week is a changeover week, the perturbation is added to the production quantity of the item in that week.
- **Max week variant:** In the max week uncertainty variant, the artificial delta perturbation is targeted to the week identified as operationally most vulnerable, according to the maximum changeover week criterion. This concentrates the uncertainty at a structurally critical point in the schedule. By applying the perturbation only in the final changeover week of each item, this variant explores an even more targeted application of the uncertainty. To create this scenario, a sudden positive perturbation is added to the production quantity of the next simulation week when this week is the maximum possible changeover week. Conceptually, this defines a worst-case stress scenario. Figure 5.4 illustrates how this perturbation is applied to the production plan. For a given material, the changeover must happen between week 44 and 46. When the next simulation week is the

maximum changeover week, the perturbation is added to the production quantity of the item in that week.

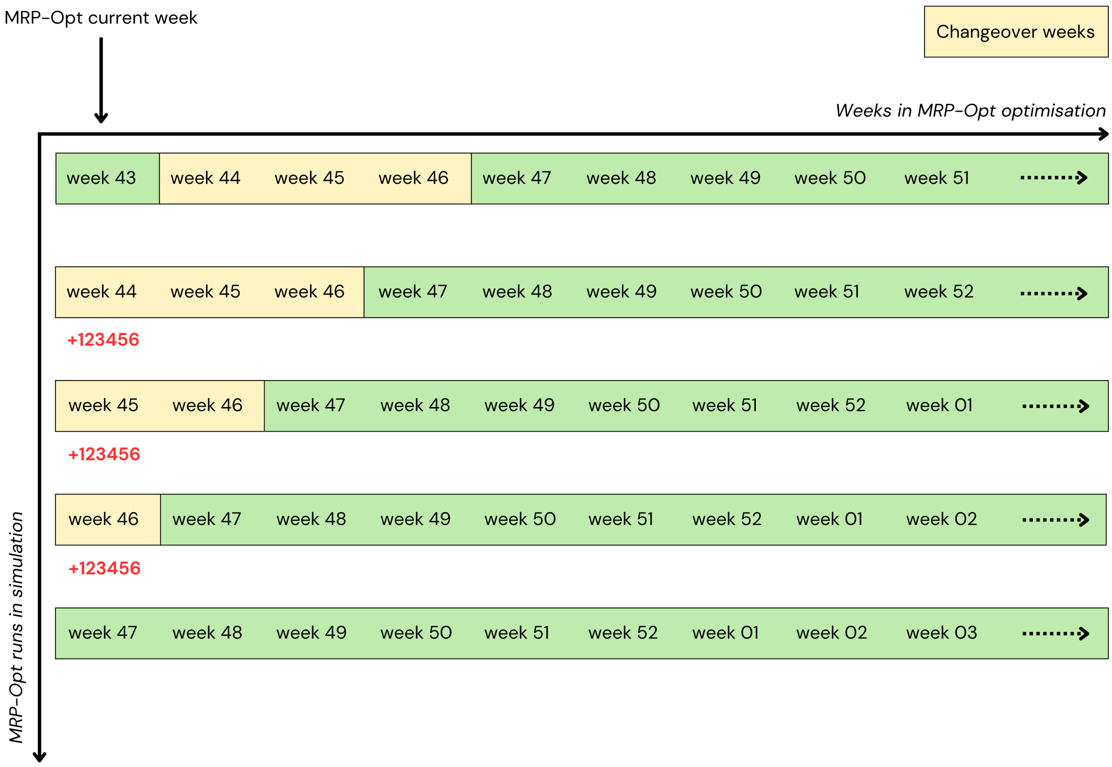

Figure 5.3: Perturbation in next simulation week when it is a changeover week

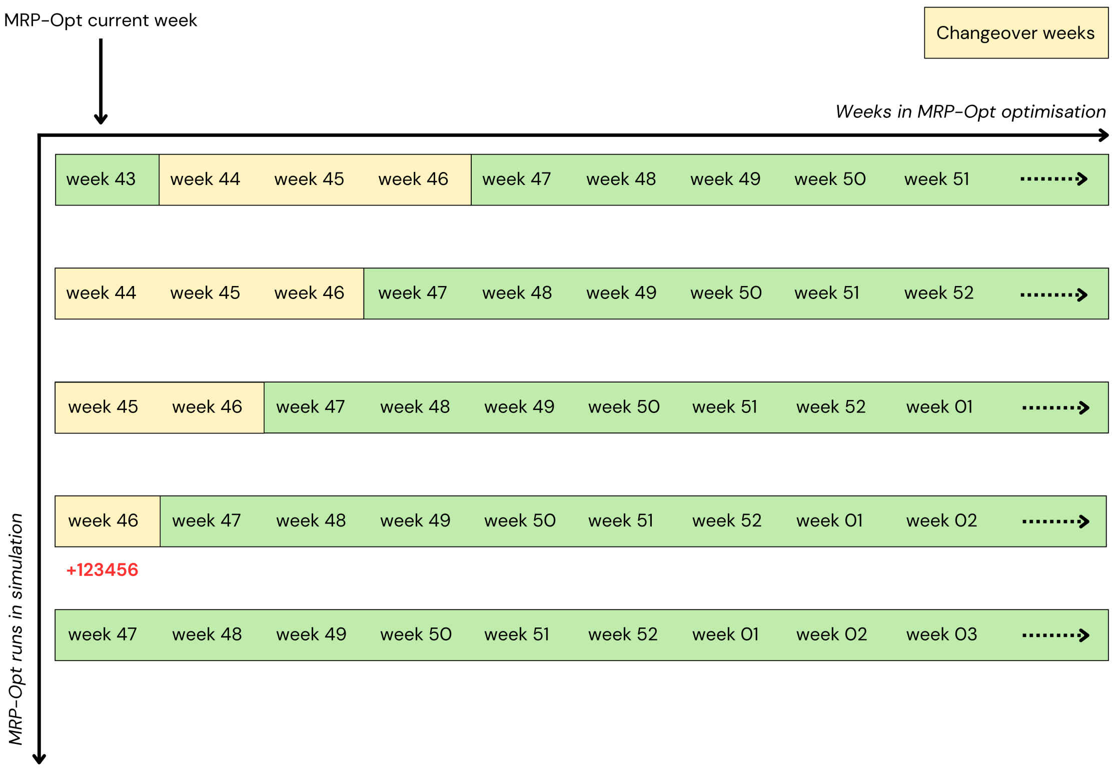

Figure 5.4: Perturbation in next simulation week when it is the maximum changeover week

5.8. Sim-Opt procedure settings

This section describes the settings that can be played with when running the Sim-Opt model. Table 5.3 gives an example, where the procedure uses a full enumeration technique to optimise the safety stock parameters of 200 materials. It uses 5 iterations for this, including 8 simulations each. In each simulation, the MRP-Opt is solved 10 times over a simulation horizon of 10 weeks, starting in week 1 of 2026.

Table 5.3: Sim-Opt procedure settings example

Optimisation method	Iterations	Simulations	Weeks simulated	Materials	Start week
Full enumeration	5	8	10	200	2026-01

These settings are of great influence on the outcome of the model and therefore important to explain in detail:

- **Optimisation method:** Determines how the solution space is searched. This influences the solution but also how fast the model converges to an optimal solution.
- **Number of Sim-Opt iterations:** Determines how many learning cycles the optimisation algorithm receives to learn the optimisation variables. The number of iterations can be set manually but it can also be linked to the optimisation method. Having more iterations generally provides more accurate solutions, but obviously also takes more time to run.
- **Number of simulations in a Sim-Opt iteration:** Determines how many simulations are run with different stochasticity, which provide a solution. The objective of an iteration is the average of all simulation solutions within that iteration. Running more simulations provides a more reliable solution for that iteration, as different kinds of stochasticity are tested with. The number of simulations can be set manually. Increasing the number of simulations significantly increases the computation time of the model, as adding one simulation means every iteration has to do one extra.
- **Number of weeks in the simulation horizon:** Determines how many weekly MRP-Opt optimisations are sequentially run within a simulation. Running multiple weekly MRP-Opt optimisations in a row is necessary as the production plans are manipulated every week, imitating the stochastic MRP process. The number of weeks can be set manually. Increasing the number of weeks in the horizon significantly increases the computation time, as adding one week means every simulation within every iteration has to do one extra.
- **Number of materials:** Determines how many materials are randomly picked and provided to the MRP-Opt to optimise for. The set of materials from which can be picked are materials which are the 'old' material in at least one changeover and are still active materials in the manufacturer's supply chain. The number of materials can be set manually. Increasing the number of materials increases the computation time as the MRP-Opt takes longer to optimise.
- **Start simulation week:** Determines which historical dataset is used for the Sim-Opt procedure. The chosen dataset provides the original production plan on which the uncertainty is added. The start week can be picked randomly by the model but can also be set manually.

6 Experimental results

This chapter presents several experiments in which the Sim-Opt framework is tested with different variants and settings. Each experiment serves a certain purpose and their results help substantiate the conclusions and discussions of this study. First, some general Sim-Opt framework settings are explored in Section 6.1. In Section 6.2, several variants of assigning uncertainty and applying HSS to the MRP-Opt model are explored. Section 6.3 tests uncertainty with different values of the artificial perturbation in the production plan. Finally, Section 6.4 shows a final experiment which zooms into weekly item-specific results in a worst-case stress scenario.

6.1. Sim-Opt settings

First off, two experiments are performed regarding the settings of the Sim-Opt framework. Experiment S1 determines how many weeks should be simulated in each simulation, while Experiment S2 discovers how many simulations should be run to provide a reliable mean objective value.

6.1.1. S1: Weeks in horizon

The goal of Experiment S1 is to determine if the number of weeks simulated in each simulation run has a large influence on the results of the objective. To do this, this experiment was performed with the settings shown in Table 6.1. This experiment does not optimise any safety stock parameters yet, as it only tests the number of weeks in each simulation. Therefore, no optimisation method was used and this experiment can be seen as 1 Sim-Opt iteration. As a reminder, Section 1.6 shows the Sim-Opt framework and describes the hierarchy between the iterations and simulations. For testing purposes, 30 materials were randomly selected from the set of active changeover materials, as mentioned in Section 5.8. 5 simulations were run to account for different levels of uncertainty. Each simulation is run over a horizon of 10 consecutive weeks, starting in week 2025-43.

Table 6.1: S1 Experiment settings

Optimisation method	Iterations	Simulations	Weeks simulated	Items	Start week
–	1	5	10	30	2025-43

Table 6.2 shows the results of the objective values for the 5 simulations in Experiment S1. To see if there are major differences between simulations, the sample mean is calculated and the *lowest* and **highest** objective value are made visible. Each simulation shows a similar increasing objective value as the weeks go by. Running more weeks, means running more MRP-Opt optimisations for which stochasticity is added to the production plan. Uncertainty is added to these production plans by adding a delta to the production levels of the previous production plan. Section 5.7 described this process in detail. Because of the way the stochasticity is added to the production plans each week, running more weeks means having more instances of adding uncertainty to the Sim-Opt procedure. This does have an effect on how realistic a simulation is. As long as at least a 5 weeks are run in a simulation, and each simulation runs the same number of weeks, this should be realistic enough.

Section 5.3 explained how the objective values are extracted from the results. The objective value of a simulation is the objective value of the final simulated week. The objective value of a Sim-Opt iteration is the average of the simulations' objective values. Although the range between the lowest and highest objective value in Table 6.2 is quite large in some weeks, the effect of this is mitigated by the fact that the

Sim-Opt objective is extracted as the mean of these objectives. Therefore, a large enough number of simulations would diminish this effect even more. So, if enough simulations are run, it does not matter how many weeks you run. The number of weeks in a simulation becomes a deterministic decision, which should be kept at the same value for consistency reasons when comparing different Sim-Opt runs.

Table 6.2: S1 Experiment result

Week [–]	Sim 1 [€]	Sim 2 [€]	Sim 3 [€]	Sim 4 [€]	Sim 5 [€]	Sample mean [€]
2025-43	**76 501**	73 555	*70 204*	74 082	73 265	73 521
2025-44	**74 775**	73 868	68 303	69 042	*67 999*	70 797
2025-45	128 014	**132 174**	129 251	129 845	*126 455*	129 148
2025-46	133 650	**141 766**	134 549	134 251	*129 586*	134 760
2025-47	152 208	123 204	150 972	*121 687*	**251 060**	159 826
2025-48	172 609	*137 532*	174 310	143 038	**262 011**	177 900
2025-49	235 760	*199 704*	240 278	227 608	**478 812**	276 432
2025-50	*318 917*	356 457	331 309	**628 665**	520 241	431 118
2025-51	630 295	644 355	674 590	**798 867**	*539 000*	657 421
2025-52	740 255	724 576	760 473	**798 981**	*719 758*	748 809

6.1.2. S2: Number of simulations

The goal of Experiment S2 is to determine what number of simulations need to be run to provide a reliable Sim-Opt iteration result. As mentioned in Experiment S1 in Section 6.1.1, the number of simulations is an important setting of the Sim-Opt procedure. Each simulation uses unique production plans for each MRP-Opt weekly optimisation. These production plans are altered based on the assignment of uncertainty. Each simulation represents a unique realisation of reality and the stochasticity influences the objective value. Therefore it is important to run multiple simulations and take the average objective value from these simulations as the Sim-Opt iteration objective value. Table 6.3 shows the settings of Experiment S2. To test how many simulations are necessary, this experiment ran 30 simulations, each with 5 weeks in the simulation horizon. The same 30 materials were sampled as in Experiment S1.

Table 6.3: S2 Experiment settings

Optimisation method	Iterations	Simulations	Weeks simulated	Items	Start week
–	1	30	5	30	2025-43

The full results of Experiment S2 can be found in Appendix G. To be able to get a reasonable conclusion from this experiment, 5, 10, 15 and 30 simulations are compared. The results of the objective values of these simulations are shown in table G.1. To compare the results, the running sample mean is calculated and added to the table. Also the sample standard deviation of the objective value of each simulation is calculated, using the same running principle. It measures the spread of the individual simulation results. The relative error expresses the variability of the mean relative to its magnitude. The formulas of these statistical metrics can be found in Appendix C.

Table 6.4: S2 Experiment result

Simulation [–]	Obsolete costs [€]	Sample mean [€]	Sample standard deviation [€]	Relative error [-]
5	179 948	163 965	21 888	0.117
10	155 190	160 879	16 921	0.065
15	160 505	158 362	18 015	0.058
30	156 543	158 679	26 289	0.059

Comparing the number of simulations in Table G.1 indicates that the sample mean at 5 simulations is already in the same order of magnitude as the sample mean at 30 simulations. The more simulations are run, the more the sample mean stabilises towards the sample mean of 30 simulations. Figure 6.1 shows the simulation objective values of all 30 simulations, together with the running sample mean. It clearly shows that the average objective value stabilises quite quickly. Although the sample standard deviation is much higher at 30 than at 10 or 15 simulations, this is due to an outlier which had an objective value of 261 491 in simulation 18. The relative error of the sample mean also stabilises around 0.060 quite early, making 10-15 simulations a reliable option. The fact that each simulation run extra results in much more computational efforts, provides a trade-off. This trade-off is between the variability of the results and the feasibility of running the Sim-Opt procedure. Therefore, the conclusion of this experiment is that the average objective value of 10 simulations provides a reliable result for one Sim-Opt iteration, while staying feasible.

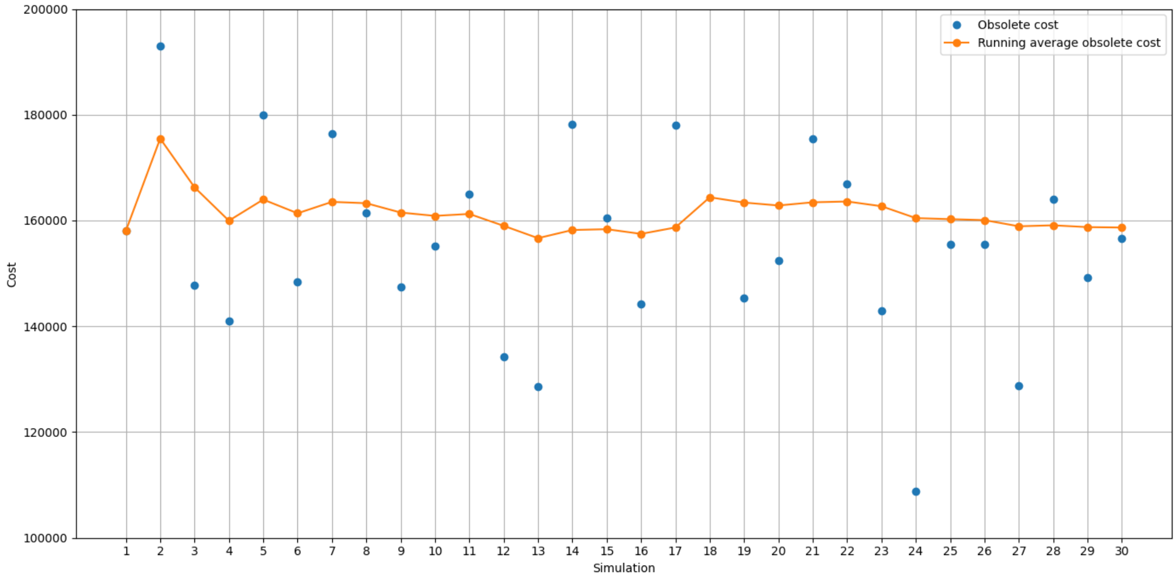

Figure 6.1: S2 Experiment simulation objective values and their running sample mean

6.2. HSS and uncertainty variants experiments

As mentioned in Sections 5.2.1 and 5.7, several variants of assigning HSS and uncertainty to the MRP-Opt are built into the Sim-Opt framework. To see which combination of variants is most useful to find improvements in performance, they are all tested. Experiments A1-A4 can be seen as a baseline to compare the uncertainty variants with, since no uncertainty is applied to the production plans in these experiments. This allows for a clear comparison in objective function value between the uncertainty assignment variants. Meanwhile, Experiments A1-D1 provide another baseline, while no HSS is assigned in these experiments. Therefore, the effect of the different types of HSS is made visible. Figure 6.2 shows an overview of the HSS and uncertainty assignment variants. For an explanation about what these variants entail, see Sections 5.2.1 and 5.7. The content of the table shows indicates what the experiments are called. The following subsections will explore the results of some of these experiments in more detail. The results of all experiments can be found in Appendices H to K.

			Uncertainty			
			A	B	C	D
			No uncertainty	Realistic uncertainty based on PDFs	Perturbation in next simulation week when it is a changeover week	Perturbation in next simulation week when it's the max changeover week
HSS	1	No HSS	A1	B1	C1	D1
	2	Constant HSS	A2	B2	C2	D2
	3	HSS in changeover weeks	A3	B3	C3	D3
	4	HSS in max changeover week	A4	B4	C4	D4

Figure 6.2: HSS and uncertainty variants experiments

To compare the objective function value results between the experiments, the Sim-Opt settings must be identical throughout. Table 6.5 shows these general Sim-Opt settings. The decision to use 10 simulations when uncertainty is applied is based on the result of Experiment S2 in Section 6.1.2. Similarly, the decision to simulate 5 weeks is a result of Experiment S1 in Section 6.1.1. As these experiments are the first exploration of parametrising the HSS, the decision was made to test HSS values of up to 7 days of sale. As full enumeration is used as an optimisation technique, this results in 7 iterations necessary for the HSS optimisation in the Sim-Opt framework. Depending on the experiments, other relevant settings are discussed throughout the following sections.

The experiments are performed using items from the MG2 material group. As mentioned in Section 5.5, materials in this group show most potential to be benefit from safety stock. The MG2 material group consists of 100 items. Initial tests resulted in MRP-Opt infeasibility, caused by some materials in this

Table 6.5: Variants experiments Sim-Opt settings

Optimisation method	Iterations	Simulations	Weeks simulated	Items	Start week
Full enumeration	7	10	5	22	2025-43

subset. After consulting a company expert and several debugging efforts, the cause of this seemed inexplicable. Similarly, materials which were used for multiple products at a location didn't provide any obsolete material costs, even when very large uncertainty was applied. These were materials were left out of the test set, resulting in 22 items to test with in the variants experiments.

6.2.1. A: No uncertainty

The A experiments are performed without uncertainty in the production plans. This means that in each MRP-Opt run, the same production plan is used. This creates a situation in which the MRP-Opt model knows exactly what the production quantity is going to be. The only thing that changes, is the week from which the MRP-Opt starts the optimisation. All results of the A experiments can be found in Appendix H. Experiment A1 specifically can be seen as a the baseline experiment, as no HSS is assigned and no uncertainty is present in the production plan. The MRP-Opt model knows exactly what is going to happen in the following weeks as the same production plan is used throughout the simulation. The results of Experiment A1 are in table 6.6. As no uncertainty is applied and no HSS parameter has to be optimised, this experiment only needs 1 iteration and 1 simulation, resulting in an obsolete material cost of €3 180.

Table 6.6: Experiment **A1** result

Iteration.simulation [–]	**Obsolete material costs** [€]	**Obsolete material stock** [pieces]	**HSS** [DOS]
1.1	3 180	15 735	0

6.2.2. B: Realistic uncertainty

The B experiments are performed with realistic uncertainty applied to the production plan. The assignment of realistic uncertainty is based on an extensive analysis of historical production plans in Section 4.2. In the realistic uncertainty variant, stochastic perturbations are generated from empirically derived PDFs and applied to 4-week blocks for all product-locations combinations. The production plan evolves under data-driven variability rather than deterministic assumptions. At the first simulated week, uncertainty is injected into the original production plan; in subsequent weeks, uncertainty is propagated on the already updated plan, which creates a path-dependent stochastic trajectory. Methodologically, this variant is the most representative for real-world operations. How this uncertainty is applied to the production plan is explained further in detail in Section 5.7. All results of the B experiments can be found in Appendix I.

Table 6.7: Experiment **B2** result

Iteration.simulation [–]	**Obsolete material costs** [€]	**Obsolete material stock** [pieces]	**HSS** [DOS]
1.M	**92 328**	**–**	**1**
2.M	113 560	–	2
3.M	136 075	–	3
4.M	156 325	–	4
5.M	178 380	–	5
6.M	199 773	–	6
7.M	222 000	–	7

One of the B experiments is Experiment B2, in which realistic uncertainty is combined with constant HSS. The results of Experiment B2 are presented in Table 6.7. The table shows the results of the 7 iterations, each representing the HSS value of the equivalent number of days of sales. In each iteration, 10 simulations were run to account for the random stochasticity applied to the production plan with each

MRP-Opt run. The mean objective value of these simulations is calculated and taken as the objective value for the iteration, as described in Section 5.3. The results show that the lowest HSS parameter results in the best objective function value, €92 328, for this experiment. Since the objective value of Experiment B1 is €70 044, the assignment of constant HSS does not provide an improvement over having no HSS in the current experimental setup. Neither do the other B experiments which test other HSS assignment variants.

6.2.3. C: Perturbation in next simulation week when it is a changeover week

To test the parametrisation of HSS in a more artificial setting, the C experiments explore the assignment of uncertainty in the form of deterministic perturbations. The perturbation is applied as a sudden delta in production volume. It is assigned to the production quantity of an item in the next simulation week, when that week is one of the item's changeover weeks. This creates an artificial bias, where the production value of next week is higher than expected. Section 5.7.1 explains how the perturbations are applied to the production plan in more detail. As a starting point, a very high perturbation of 123456 is chosen to see its effect on the objective function values. Because the uncertainty is deterministically chosen, there is no need to run multiple simulations in these experiments. All results of the C experiments can be found in Appendix J.

One of these experiments, is experiment C3. This experiment explores the artificial perturbation in combination with HSS that is applied to the changeover weeks of each item individually only. This way of assigning HSS to the MRP-Opt model makes sure HSS is maintained during the weeks in which it could be of most use. The results of Experiment C3 are in table 6.8. The results show that the HSS parameter value of 1 provides the lowest objective function value (€23 948), similar to the other variants experiments. A steady increase in objective value can be seen, the more HSS is assigned to the MRP-Opt model. Experiment C1 (no HSS) gives an objective of €2 920, meaning that the assignment of HSS during changeover weeks does not provide a better objective value.

Table 6.8: Experiment **C3** result

Iteration (.simulation) [–]	**Obsolete material costs** [€]	**Obsolete material stock** [pieces]	**HSS** [DOS]
1.1	**23 948**	**305 262**	**1**
2.1	44 310	594 365	2
3.1	65 338	886 168	3
4.1	85 700	1 175 271	4
5.1	106 125	1 465 035	5
6.1	127 089	1 756 177	6
7.1	147 486	2 045 643	7

6.2.4. D: Perturbation in next simulation week when it's the maximum changeover week

Similarly to the experiments in Section 6.2.3, the D experiments apply artificial uncertainty in the form of deterministic perturbations to the production plans. However, the HSS is now only applied for each item when the next simulation week is the maximum changeover week. This provides an exact week in which the perturbation is assigned for each item, creating a very specific experimental environment to test the HSS assignments in. Section 5.7.1 explains how the perturbations are applied to the production plan in more detail. A very high perturbation of 123456 is chosen as a starting point to see its effect on the objective. Because the uncertainty is deterministically chosen, there is no need to run multiple simulations in these experiments. All results of the D experiments can be found in Appendix K.

Experiment D4 combines this artificial uncertainty with HSS assignment solely in the maximum changeover week. The results of Experiment D4 are in table 6.9. Since both are assigned to the same exact week, this experiment showed most potential for HSS to be beneficial. However, it does not outperform the objective value of €2 889 for experiment D1, without HSS. The same pattern as in all other variants experiments where HSS was assigned can be seen in the table, where the objective value simply increases, the more HSS is assigned to the MRP-Opt model.

Table 6.9: Experiment **D4** result

Iteration (.simulation) [–]	**Obsolete material costs** [€]	**Obsolete material stock** [pieces]	**HSS** [DOS]
1.1	**24 502**	**308 395**	**1**
2.1	45 387	600 271	2
3.1	67 000	895 567	3
4.1	87 884	1 187 443	4
5.1	108 799	1 479 679	5
6.1	130 382	1 774 615	6
7.1	151 297	2 066 851	7

6.2.5. HSS and uncertainty variants results

The goal of the HSS and uncertainty variants experiments is to compare them based on their best objective function values. Meanwhile, the experiments also gave an insight into the usefulness of HSS in the MRP-Opt model under the current experimental setup. The objective values of Experiments A1-D4 are shown in table 6.3. This table provides a visible comparison between the HSS and uncertainty variants. The complete results of Experiments A to D are presented in Appendices H to K.

			Uncertainty			
			A	B	C	D
			No uncertainty	Realistic uncertainty based on PDFs	Perturbation in next simulation week when it is a changeover week	Perturbation in next simulation week when it is the max changeover week
HSS	1	No HSS	3 180	70 044	2 920	2 889
	2	Constant HSS	28 612	92 328	28 240	28 209
	3	HSS in changeover weeks	24 321	88 946	23 948	23 918
	4	HSS in max changeover week	24 905	79 467	24 471	24 502

Figure 6.3: Variants experiments results

A first observation is that the introduction of realistic uncertainty substantially increases the objective function value compared to the no-uncertainty baseline. In Experiment A1, where no uncertainty and no HSS are applied, the obsolete material costs amount to only €3 180. In contrast, the corresponding no-HSS experiment realistic uncertainty result in a substantially higher objective value of €70 044 in B1. This confirms that the realistic uncertainty introduces relevant stress into the MRP-Opt environment. The major increase shows that deviations in production quantities strongly affect obsolete material costs in a rolling-horizon changeover setting. Meanwhile, the introduction of an artificial perturbation in the production plan resulted in a decrease in objective function value, for both variants. These perturbations resulted in values of €2 920 in C1 and €2 889 in D1. This is not as expected, since large spikes in production value should cause material procurements with MOQ. These could result in obsolete material costs as well. It is possible that the current perturbation value of 123 456 is too large, causing current obsolete materials to be used anyway. However, the expectation would be that new orders are planned as well so that the large production volume can be produced in full. Another explanation could be that the timing of the perturbation is causing the MRP-Opt to not have enough time to order stock for the increase in production, causing the obsolete stock not to increase either.

The no-uncertainty experiments also illustrate the direct effect of adding HSS when no stochasticity is present. In Experiments A2 to A4, the MRP-Opt model uses the same production plan throughout the simulation and therefore has perfect information about future production quantities. Under these deterministic conditions, HSS does not provide protection against demand fluctuations in the existing production plan. Instead, it only forces the model to maintain additional inventory that is not required by the production plan. This results in higher obsolete material costs compared to A1. For example, the HSS setting in A3 shows an increase of the objective value to €24 321, compared to €3 180 without HSS. This result is expected and confirms that HSS creates additional obsolete stock costs when the buffer is not needed.

The same general pattern is observed when uncertainty is introduced. For all uncertainty variants experiments, the best-performing HSS value is consistently the lowest tested value of 1 day of sales (DOS). But they never outperform the baseline of no HSS. In Experiment B2, where realistic uncertainty is combined with constant HSS, the best objective value is €92 328 at 1 DOS. This is still worse than the no-HSS realistic uncertainty baseline of €70 044 in Experiment B1. Experiment B4, with HSS in the maximum changeover week, performs better than the HSS in changeover weeks variant B3. This is striking, as in A, C and D, the HSS in max changeover week performs slightly worse than the HSS in all changeover weeks. This may indicate that under realistic uncertainty conditions, the more targeted HSS assignment might be a better pick. The artificial uncertainty experiments showed that the HSS never outperformed the baseline. Experiment D4 was expected to provide one of the most favourable settings for HSS, because both the perturbation and the HSS are concentrated in the same changeover week. Nevertheless, the best objective value in D4 is €24 502, which remains higher than the no-HSS baseline of €2 889 in D1.

These results indicate that, within the current experimental setup, adding HSS to the MRP-Opt model does not improve the objective function value compared to the baseline without HSS. This conclusion holds for the no-uncertainty setting, the realistic uncertainty setting and the artificial uncertainty settings. The monotonic increase in objective value as the HSS parameter increases further supports this conclusion. Across the experiments, higher HSS values lead to larger obsolete material stocks and therefore to higher obsolete material costs. This suggests that the additional inventory forced by the HSS constraint is not sufficiently used to compensate for uncertainty before it becomes obsolete. In other words, the protective value of the buffer is not reflected in the current objective function values, while the negative effect of holding additional changeover material is directly penalised through obsolete stock costs.

Although MG2 was selected because lumpy materials are expected to have relatively high potential for safety stock, the tested set may still contain items for which HSS is not beneficial. If only a few items benefit from the buffer while the majority experience higher obsolescence, the aggregated objective value will still deteriorate. To test this theory, it would be worthwhile to zoom into material-specific obsolete stock costs to see what happens on a item-level. Since only 5 weeks were simulated, not all changeover weeks of every item fell into this simulation horizon. This could have a negative effect on the artificial uncertainty experiments, as HSS might have been enforced, without the perturbation being added to the production plan. Therefore, the simulation horizon could be broadened to ensure all changeover periods are included.

The artificial uncertainty experiments remain useful, despite not producing an improvement over the baseline. Their purpose was to create controlled stress scenarios in which HSS should theoretically have a stronger chance of being beneficial. In the C experiments, the perturbation is applied when the next simulation week is a changeover week. In the D experiments, the perturbation is concentrated even more specifically in the maximum changeover week. These settings help isolate the relationship between perturbation timing and HSS timing. The fact that HSS still does not improve the objective value under these favourable conditions suggests that the current HSS formulation and objective function may not be sufficient to create measurable improvements in obsolete stock costs. For the perturbation, a large value of 123456 was deterministically chosen to test with in these experiments. Therefore, it could be useful to test other perturbation sizes to see if these provide conditions under which the current HSS formulation performs well.

Overall, the HSS and uncertainty variant experiments are useful because they separate the effects of uncertainty assignment and HSS timing. The A experiments confirm the deterministic baseline effect of HSS. The B experiments show how HSS behaves under data-driven demand variability. The C and D experiments test whether HSS becomes beneficial under more targeted artificial disruptions around changeover moments. Together, these experiments provide a structured comparison of the available modelling variants and show that, within the current Sim-Opt setup, HSS does not outperform the no-HSS baseline.

6.3. Perturbation size experiments

The experiments in Section 6.2 showed that HSS did not improve the aggregated objective value for the tested MG2 item set. The experiments using an artificial perturbation in the production plan gave lower objective values, which is counter intuitive. That is why this section explores different perturbation sizes. Meanwhile, the experiments in Section 6.2 were evaluated over a set of multiple items. This means that possible improvements for individual items could have been offset by worse results for other items. Therefore, the perturbation experiments are only tested with 5 items, to analyse whether HSS can be beneficial for specific material-location-product combinations. The goal of these experiments is not to find a general HSS policy for the full material group, but to investigate the behaviour of HSS under more controlled conditions. By testing only a small number of individual items, the effect of the HSS parameter can be studied in more detail. This makes it possible to see whether the limited performance found in Section 6.2 is caused by HSS being generally ineffective, or whether the potential benefit is more item-specific.

The perturbation size experiments use the artificial uncertainty in the maximum changeover week variant from Section 6.2. This variant is most relevant for this analysis, because it creates situations in which HSS should theoretically have the highest chance of being useful. These experiments combine this uncertainty variant with the assignment of HSS to maximum changeover weeks only. A positive perturbation applied specifically in the maximum changeover week, together with the HSS assignment in the same week, creates a controlled stress scenario in which a safety stock buffer could help prevent unfavourable planning outcomes.

A limitation of the previous experiments was that not all changeover weeks of the tested items occurred within the simulated horizon. As a result, HSS could be enforced without the corresponding demand perturbation taking place during the simulation. To avoid this, five items were selected for which the relevant changeover weeks fall within the expanded simulated period of 11 weeks. They were selected so that the test set is a varied sample based on the item's characteristics. Table 6.10 shows these individually tested items. The table includes the changeover weeks, piece cost, MOQ and one-day-of-sales for each item. These characteristics are important because they determine how large the absolute HSS buffer becomes and how costly additional obsolete stock can become.

Table 6.10: Perturbation size Experiments test items

Item [–]	Changeover weeks [–]	Piece cost [€]	MOQ [pieces]	1 day of sales [pieces]
1	2025-47; 2025-49	0.085176	5 760	48 703
2	2025-46; 2025-46	0.11044	5 446	9 876
3	2025-47; 2025-51	0.0064976	370 000	58 178
4	2025-46; 2025-52	0.06519	24 000	10 084
5	2025-52; 2025-52	0.2470703	2 700	4 004

The Sim-Opt settings for the perturbation size experiments are shown in Table 6.11. Compared to the experiments in Section 6.2, the number of simulated weeks is increased from 5 to 11. This ensures that all selected item changeover weeks are included in the simulation horizon. Full enumeration is again used to test HSS values from 1 to 7 days of sales. By using fewer items but a longer simulation horizon, the experiments becomes more suitable for analysing the timing of the perturbation and the HSS assignment.

Table 6.11: Perturbation size Experiments Sim-Opt settings

Optimisation method	Iterations	Simulations	Weeks simulated	Items	Start week
Full enumeration	7	10	11	5	2025-43

Figure 6.4 shows all the results of the perturbation size experiments. Appendix L shows the detailed results of all experiments in this figure.

			Uncertainty				
			A	D			
			No uncertainty	Perturbation in next simulation week when it is the max changeover week			
				P = 123	P = 1 234	P = 12 345	P = 123 456
HSS	1	No HSS	575	596	418	0	0
	4	HSS in max changeover week	8 327	7 507	7 360	7 264	7 264

Figure 6.4: Perturbation size experiments results

The perturbation size experiments show that the effect of the artificial demand spike is not monotonic in the no-HSS case. In D1, where the perturbation is applied in the maximum changeover week but no HSS is assigned, the smallest perturbation $P = 123$ results in slightly higher obsolete costs (€596) than the no-uncertainty baseline A1 (€575). However, increasing the perturbation to $P = 1234$ reduces obsolete costs to €418, while the two largest perturbations, $P = 12345$ and $P = 123456$, both result in zero obsolete stock and therefore zero costs. This is counter-intuitive, because a larger last-minute demand increase would be expected to increase the probability of additional procurement and MOQ-related leftover stock. Instead, the results suggest that the larger perturbations consume the stock that would otherwise have become obsolete, without triggering additional obsolete material in the final evaluated week. Possibly, the timing of the perturbation is causing the MRP-Opt to not have enough time to order stock for the increase in production, causing the obsolete stock to not increase either.

When HSS is assigned only in the maximum changeover week, as in D4, the results change substantially. For every perturbation size, the introduction of HSS immediately creates a much higher obsolete stock level than the no-HSS variant. With $P = 123$, the total obsolete costs increase from €596 in D1 to €7 507 at a HSS of 1 day of sales (DOS). For $P = 1234$, this increase is from €418 to €7 360. For $P = 12345$ and $P = 123456$, the difference is even clearer: the no-HSS variant has zero obsolete costs, whereas D4 with 1 DOS results in €7 264 of obsolete costs. Therefore, the HSS buffer does not reduce obsolete stock in any perturbation-size setting. Instead, it creates additional leftover inventory that is penalised after the changeover.

Comparing the different perturbation sizes within D4 shows that larger perturbations slightly reduce the obsolete costs caused by HSS, rather than increasing them. At 1 DOS, the obsolete costs decrease from €7 507 for $P = 123$, to €7 360 for $P = 1234$, and to €7 264 for both $P = 12345$ and $P = 123456$. A similar pattern can be seen at higher HSS values: at 7 DOS, the costs are €51 188 for $P = 123$, €51 010 for $P = 1234$, and €50 846 for both larger perturbations. This indicates that the positive perturbation partly consumes the HSS buffer before the changeover, thereby slightly reducing the amount of stock left obsolete. However, this reduction is small compared to the additional obsolete stock created by enforcing HSS in the first place.

The item-level results help explain this behaviour. In D4 with $P = 123$ and 1 DOS, the obsolete stock of items 1 to 4 is almost exactly equal to their one-day-of-sales value (Table 6.10: 48 703, 9 876, 58 178 and 10 084 pieces respectively. This shows that the HSS constraint is directly visible in the obsolete stock outcome. The model is forced to retain at least the HSS quantity in the maximum changeover week, and this quantity becomes obsolete after the changeover. Item 5 deviates somewhat from this exact pattern, with 4 990 obsolete pieces compared to 4 004 pieces for 1 DOS, suggesting that other planning logic, existing stock, or ordering effects also influence the final obsolete quantity for this item.

The comparison between $P = 12345$ and $P = 123456$ is especially interesting. Both perturbation sizes produce identical results in D1 and D4. In D1, both lead to zero obsolete costs, and in D4 all HSS levels from 1 to 7 DOS give exactly the same obsolete stock and cost values for both perturbation sizes. This suggests that, beyond a certain perturbation threshold, the MRP-Opt response no longer changes. A likely explanation is that the perturbation pushes the model into the same planning decision structure, after which increasing the production quantity further does not affect the final obsolete stock outcome. This suggests that there is either some fundamental reason in the current setup of the Sim-Opt framework or the MRP-Opt model that causes the HSS to not be beneficial, or that the intuitive expectation that safety stock can improve performance with high perturbations is unrealistic.

Overall, the perturbation size experiments do not show a case in which HSS improves the objective value. The no-HSS variants always perform better than the corresponding HSS-in-maximum-changeover-week variants. However, the results do show that perturbation size affects how much of the HSS buffer remains obsolete. Smaller perturbations leave more of the buffer unused, while larger perturbations consume part of it. This means that the intuition behind HSS is partly visible, since additional demand can indeed reduce leftover stock. Nevertheless, the tested HSS quantities are too large relative to the benefit created by the perturbation, so the net effect remains negative. This supports the need to test lower HSS values, especially because the one-day-of-sales value already represents a large absolute quantity for some of the selected items.

6.4. Proof-of-concept experiment

Based on the experiments in Sections 6.2 and 6.3, this final proof-of-concept is proposed to find a scenario in which HSS might be beneficial. To create a worst-case stress scenario, the artificial uncertainty variant with **perturbation in the maximum changeover week** is used. It focuses on one item only, to see how it responds to different HSS and perturbation settings. The results of the previous experiments all showed that 1 DOS performed best, without outperforming no HSS. If zero DOS would have been included in the HSS experiments, it would have been chosen. Therefore, it is worthwhile testing lower HSS values in combination with the different perturbation sizes to see how the objective function value reacts to this. Table 6.12 shows the item that is chosen to test with in this experiment. This experiment will test only with the **HSS in maximum changeover week** variant. Therefore, this item was chosen because its minimum and maximum changeover week are the same week. It is expected that the impact of the perturbations and HSS assignment would be most easily visible when an item has only 1 changeover week. Similarly, its MOQ and one-day-of-sales values are in approximately the same order of magnitude as some of the perturbation sizes.

Table 6.12: Proof-of-concept Experiment test item

Item [–]	**Changeover weeks** [–]	**Piece cost** [€]	**MOQ** [pieces]	**1 day of sales** [pieces]
5	2025-52; 2025-52	0.2470703	2 700	4 004

The Sim-Opt settings in Table 6.13 will be used for this experiment. Here, 17 iterations are used because this experiment will test the HSS parameters 0.1 to 0.9 with steps of 0.1, and 1 to 7 with steps of 1. Since this experiment only focuses on 1 item, and its changeover must take place in week 2025-52, the start week of this experiment is set in week 2025-50 and the simulation will run for 5 weeks. This should capture the effects of HSS and uncertainty around the changeover week, while keeping the experiment tractable.

Table 6.13: Proof-of-concept Experiment Sim-Opt settings

Optimisation method	Iterations	Simulations	Weeks simulated	Items	Start week
Full enumeration	17	1	5	1	2025-50

The results in Tables 6.14 and 6.15 show that the behaviour of the proof-of-concept experiment is strongly dependent on both the perturbation size and the imposed HSS level. Since week 2026-02 is the final week in the five-week simulation horizon, the obsolete stock values in this week are used as the final result of each simulation. In the no-HSS case, the obsolete stock costs do not increase monotonically with the perturbation size. Without uncertainty, the final obsolete stock equals 1 151 pieces, corresponding to €284. With a small perturbation of $P = 123$, this increases to 2 290 pieces and €566, while $P = 1234$ results in 1 692 pieces and €418. However, for the larger perturbation $P = 12345$, the obsolete stock becomes zero, and for $P = 123456$, only 400 pieces remain obsolete, corresponding to €99. This indicates that a larger positive perturbation in week 2025-52 can consume more of the old material before or during the changeover, thereby reducing the remaining obsolete stock at the end of the simulation. This is in line with what was seen in the results of Section 6.3. This result is important because it shows that the perturbation does not necessarily worsen the obsolete stock objective. In this specific item-level setting, the perturbation represents additional demand in the maximum changeover week. When this additional demand is sufficiently large, it helps deplete the old material stock that would otherwise become obsolete after the changeover.

When HSS is introduced in the maximum changeover week, the results change considerably. Even low HSS values start to impose additional material stock that can remain after the changeover. For example, at 0.1 DOS, the results for no uncertainty, $P = 123$, and $P = 1234$ remain equal to the no-HSS case, but for $P = 12345$, the final obsolete stock increases from zero to 400 pieces (1/10 of the one-day-of-sales value), corresponding to €99. At 0.2 DOS, the same perturbation results in 801 obsolete pieces and €198. This shows that once the HSS constraint becomes binding, the enforced buffer is directly reflected in the obsolete stock outcome. Instead of only protecting against uncertainty, the HSS creates additional material that must still be depleted before the changeover. If the perturbation is not large enough to consume this buffer, the remaining HSS quantity becomes obsolete.

When looking at the results of 0.2 DOS and 0.3 DOS, you can see a sudden increase in obsolete material stock and costs. This is because, 0.1 and 0.2 DOS are equal to a HSS of 400 and 801 and therefore stay below the stock level (1 151) that was already obsolete in the 0 DOS scenario. When increasing the HSS to 0.3, a HSS of 1 201 is necessary, which is higher than 1 151. Therefore, more stock is necessary and you see that the obsolete material stock increases with 2 700 pieces to 3 851, which is equal to the MOQ. This indicates that the MOQs are still enforced in the MRP-Opt tool, but somehow orders only seem to be executed when more HSS is necessary, and not when a high increase in production quantity is applied.

The comparison across HSS levels shows a clear upward trend in obsolete stock costs as the safety stock level increases. For the larger HSS values, the obsolete stock increasingly resembles the imposed HSS quantity itself. At 1 DOS, the item has a safety stock of 4 004 pieces, which corresponds to €989 in obsolete material costs for $P = 12345$ and $P = 123456$. For higher HSS levels, this effect becomes even stronger: at 7 DOS, the final obsolete stock reaches 28 028 pieces and €6 925 under $P = 123456$. This pattern indicates that the hard safety stock constraint dominates the final obsolete stock result. The buffer is enforced in the model, but after the changeover the same buffer becomes leftover stock, which is penalised by the objective function.

The effect of perturbation size is therefore twofold. On the one hand, larger perturbations increase demand in the maximum changeover week and can help consume old material stock, which lowers obsolete stock costs in the no-HSS case. On the other hand, when HSS is active, the perturbation must first be large enough to consume both the original remaining stock and the additional enforced buffer. If this does not happen, the HSS quantity remains in the system and becomes obsolete. This explains why larger perturbations can slightly reduce the negative effect of HSS, but do not reverse it. The buffer is partly consumed by the demand spike, but not sufficiently to outperform the no-HSS solution. Meanwhile, the perturbations do not result in procurements with MOQs that should intuitively also result in more obsolete material costs. This is visible in the no HSS with perturbations experiments.

Compared with the no-HSS variant, the HSS in the maximum changeover week does not improve the objective value for this proof-of-concept item. The no-HSS variant consistently performs better or equal in terms of obsolete material costs, especially for the larger perturbations. Overall, the proof-of-concept experiment confirms the pattern observed in the previous perturbation size experiments: HSS can be made visible in the item-level results, but it does not improve obsolete material costs under the current Sim-Opt framework setup. The experiment does demonstrate the mechanism that was expected from the artificial perturbation setup: a positive perturbation in the maximum changeover week can reduce obsolete stock by consuming old material before it becomes obsolete. However, the tested HSS values introduce additional stock that is more costly than the benefit created by this demand increase. This suggests that, for this item in the current setup, the optimal policy remains to use no HSS. This experiment confirms the results in the previous sections on a weekly and item level.

Table 6.14: Proof-of-concept Experiment result **part 1**

HSS	**Week**	**No uncertainty**		**P = 123**		**P = 1234**		**P = 12345**		**P = 123456**	
		costs	stock	costs	stock	costs	stock	costs	stock	costs	stock
[DOS]	[–]	[€]	[pieces]	[€]	[pieces]	[€]	[pieces]	[€]	[pieces]	[€]	[pieces]
0	2025-50	284	1151	284	1151	284	1151	284	1151	284	1151
	2025-51	284	1151	284	1151	284	1151	284	1151	284	1151
	2025-52	284	1151	566	2290	418	1692	0	0	99	400
	2026-01	284	1151	566	2290	418	1692	0	0	99	400
	2026-02	**284**	**1151**	**566**	**2290**	**418**	**1692**	**0**	**0**	**99**	**400**
0.1	2025-50	284	1151	284	1151	284	1151	284	1151	284	1151
	2025-51	284	1151	284	1151	284	1151	284	1151	284	1151
	2025-52	284	1151	566	2290	418	1692	99	400	198	801
	2026-01	284	1151	566	2290	418	1692	99	400	198	801
	2026-02	284	1151	566	2290	418	1692	99	400	198	801
0.2	2025-50	284	1151	284	1151	284	1151	284	1151	951	3851
	2025-51	284	1151	284	1151	284	1151	284	1151	951	3851
	2025-52	284	1151	566	2290	418	1692	198	801	297	1201
	2026-01	284	1151	566	2290	418	1692	198	801	297	1201
	2026-02	284	1151	566	2290	418	1692	198	801	297	1201
0.3	2025-50	951	3851	951	3851	951	3851	951	3851	951	3851
	2025-51	951	3851	951	3851	951	3851	951	3851	951	3851
	2025-52	951	3851	566	2290	418	1692	297	1201	396	1602
	2026-01	951	3851	566	2290	418	1692	297	1201	396	1602
	2026-02	951	3851	566	2290	418	1692	297	1201	396	1602
0.4	2025-50	951	3851	951	3851	951	3851	951	3851	951	3851
	2025-51	951	3851	951	3851	951	3851	951	3851	951	3851
	2025-52	951	3851	566	2290	418	1692	396	1602	495	2002
	2026-01	951	3851	566	2290	418	1692	396	1602	495	2002
	2026-02	951	3851	566	2290	418	1692	396	1602	495	2002
0.5	2025-50	951	3851	951	3851	951	3851	951	3851	951	3851
	2025-51	951	3851	951	3851	951	3851	951	3851	951	3851
	2025-52	951	3851	566	2290	1085	4392	495	2002	593	2402
	2026-01	951	3851	566	2290	1085	4392	495	2002	593	2402
	2026-02	951	3851	566	2290	1085	4392	495	2002	593	2402

Table 6.15: Proof-of-concept Experiment result **part 2**

HSS	Week	No uncertainty		P = 123		P = 1234		P = 12345		P = 123456	
		costs	stock	costs	stock	costs	stock	costs	stock	costs	stock
[DOS]	[–]	[€]	[pieces]	[€]	[pieces]	[€]	[pieces]	[€]	[pieces]	[€]	[pieces]
0.6	2025-50	951	3851	951	3851	951	3851	951	3851	951	3851
	2025-51	951	3851	951	3851	951	3851	951	3851	951	3851
	2025-52	951	3851	593	2402	1085	4392	593	2402	693	2803
	2026-01	951	3851	593	2402	1085	4392	593	2402	693	2803
	2026-02	951	3851	593	2402	1085	4392	593	2402	693	2803
0.7	2025-50	951	3851	951	3851	951	3851	951	3851	951	3851
	2025-51	951	3851	951	3851	951	3851	951	3851	951	3851
	2025-52	951	3851	1233	4990	1085	4392	693	2803	791	3203
	2026-01	951	3851	1233	4990	1085	4392	693	2803	791	3203
	2026-02	951	3851	1233	4990	1085	4392	693	2803	791	3203
0.8	2025-50	951	3851	951	3851	951	3851	951	3851	951	3851
	2025-51	951	3851	951	3851	951	3851	951	3851	951	3851
	2025-52	951	3851	1233	4990	1085	4392	791	3203	890	3604
	2026-01	951	3851	1233	4990	1085	4392	791	3203	890	3604
	2026-02	951	3851	1233	4990	1085	4392	791	3203	890	3604
0.9	2025-50	951	3851	951	3851	951	3851	951	3851	1619	6551
	2025-51	951	3851	951	3851	951	3851	951	3851	1619	6551
	2025-52	951	3851	1233	4990	1085	4392	890	3604	1088	4404
	2026-01	951	3851	1233	4990	1085	4392	890	3604	1088	4404
	2026-02	951	3851	1233	4990	1085	4392	890	3604	1088	4404
1	2025-50	1619	6551	1619	6551	1619	6551	1619	6551	1619	6551
	2025-51	1619	6551	1619	6551	1619	6551	1619	6551	1619	6551
	2025-52	1619	6551	1233	4990	1085	4392	989	4004	989	4004
	2026-01	1619	6551	1233	4990	1085	4392	989	4004	989	4004
	2026-02	1619	6551	1233	4990	1085	4392	989	4004	989	4004
2	2025-50	2286	9251	2286	9251	2286	9251	2286	9251	2286	9251
	2025-51	2286	9251	2286	9251	2286	9251	2286	9251	2286	9251
	2025-52	2286	9251	2567	10390	2419	9792	1979	8008	1979	8008
	2026-01	2286	9251	2567	10390	2419	9792	1979	8008	1979	8008
	2026-02	2286	9251	2567	10390	2419	9792	1979	8008	1979	8008
3	2025-50	3620	14651	3620	14651	3620	14651	3620	14651	3620	14651
	2025-51	3620	14651	3620	14651	3620	14651	3620	14651	3620	14651
	2025-52	3620	14651	3234	13090	3086	12492	2968	12012	2968	12012
	2026-01	3620	14651	3234	13090	3086	12492	2968	12012	2968	12012
	2026-02	3620	14651	3234	13090	3086	12492	2968	12012	2968	12012
4	2025-50	4287	17351	4287	17351	4287	17351	4287	17351	4287	17351
	2025-51	4287	17351	4287	17351	4287	17351	4287	17351	4287	17351
	2025-52	4287	17351	3957	16016	4421	17892	3957	16016	3957	16016
	2026-01	4287	17351	3957	16016	4421	17892	3957	16016	3957	16016
	2026-02	4287	17351	3957	16016	4421	17892	3957	16016	3957	16016
5	2025-50	4954	20051	4954	20051	4954	20051	4954	20051	4954	20051
	2025-51	4954	20051	4954	20051	4954	20051	4954	20051	4954	20051
	2025-52	4954	20051	5235	21190	5088	20592	4946	20020	4946	20020
	2026-01	4954	20051	5235	21190	5088	20592	4946	20020	4946	20020
	2026-02	4954	20051	5235	21190	5088	20592	4946	20020	4946	20020
6	2025-50	6288	25451	6288	25451	6288	25451	6288	25451	6288	25451
	2025-51	6288	25451	6288	25451	6288	25451	6288	25451	6288	25451
	2025-52	6288	25451	5936	24024	6422	25992	5936	24024	5936	24024
	2026-01	6288	25451	5936	24024	6422	25992	5936	24024	5936	24024
	2026-02	6288	25451	5936	24024	6422	25992	5936	24024	5936	24024
7	2025-50	6955	28151	6955	28151	6955	28151	6955	28151	6955	28151
	2025-51	6955	28151	6955	28151	6955	28151	6955	28151	6955	28151
	2025-52	6955	28151	7237	29290	7089	28692	6925	28028	6925	28028
	2026-01	6955	28151	7237	29290	7089	28692	6925	28028	6925	28028
	2026-02	6955	28151	7237	29290	7089	28692	6925	28028	6925	28028

7 Discussion

This chapter discusses the findings of this research in relation to the main research question: *How can simulation-based optimisation be used to parametrise safety stocks of changeover materials in a deterministic MIP model, such that rolling horizon performance improves in a stochastic MRP environment?* Building on the literature review, the current state analysis, the data analysis, the modelling framework and the experimental results, this discussion reflects on the meaning of the obtained results. The implications for the manufacturer's MRP process and the methodological limitations of the proposed framework are also discussed.

7.1. Main interpretation of the results

The central finding of this thesis is that the proposed simulation-based optimisation (Sim-Opt) framework was technically capable of parametrising safety stock levels for changeover materials, but did not lead to improvements in objective function value. More specifically, the addition of hard safety stock continuously resulted in an objective value that was worse than the baseline. The variants experiments in Section 6.2 indicate that the addition of the optimised hard safety stock values did not reduce obsolete stock costs at a multi-item level. Similarly, Section 6.4 provided the same finding for an individual item. This is important because it provides a nuanced answer to the main research question: simulation-based optimisation can indeed be used to parametrise safety stock levels in a deterministic MIP-based MRP environment, but the resulting policy does not necessarily improve rolling-horizon performance within the current Sim-Opt framework setup.

This outcome should not be interpreted as proof that safety stock is irrelevant for changeover materials. Rather, it suggests that in the manufacturer's current MRP setting, the value of safety stock is highly dependent on how the performance is measured. The results of the proof-of-concept experiment in Section 6.4 indicate that the benefits of safety stock, if present, are not visible in the current objective function. This study highlights the difficulty of translating a robustness measure that is intuitively useful in practice into a measurable system-wide improvement in a large industry-size application.

A second important interpretation is that the thesis demonstrates a distinction between methodological feasibility and practical viability. The proposed Sim-Opt framework discussed in Section 1.6 and elaborated in Chapter 5 worked as intended from a modelling perspective. Demand uncertainty could be assigned to the production plans, the MRP-Opt could be run sequentially in a rolling-horizon manner and the Full Enumeration method could optimise safety stock parameters while treating the original MIP model as a black box. The performed experiments in Chapter 6 demonstrate that the framework is operational and capable of exploring the parameter space. However, technical feasibility alone does not guarantee that the chosen decision variable and evaluation metric will produce business value. This distinction is highly important when interpreting the results of this thesis.

7.2. Understanding the limited performance

The key question arising from the experimental results is why the optimised safety stock parameters did not outperform the baseline. The most important explanation lies in the structure of the optimisation objective chosen for the Sim-Opt framework. As described in Section 5.3, the optimisation procedure was built around the obsolete stock decision variable. The safety stock optimisation uses obsolete stock costs as the sole optimisation objective. This design choice was necessary because the MRP-Opt output did not provide sufficient information on inventory holding costs or rush-order costs. This data is needed to model the full trade-off described in Section 1.1 between inventory costs, obsolete stock costs and costs associated with last-minute orders. As a result, the Sim-Opt framework only captures one dimension of the financial consequences of safety stock. Any positive economic effect of adding safety stock such as reducing emergency replenishments or operational instability remained outside the evaluation logic of the Sim-Opt framework. The additional stock created by HSS is penalised when it becomes obsolete, while the robustness benefit it may provide remains invisible in the optimisation result.

This effect is reinforced by the specific characteristics of changeover materials. For regular materials, additional inventory can often still be consumed in later periods, but for changeover materials the usable time window is limited. Once the changeover has taken place, remaining stock of the old material directly becomes obsolete. The experiments confirm this mechanism: applying HSS forces the MRP-Opt model to retain additional stock during the changeover window. If the simulated demand increase is not large enough to consume this buffer, the remaining stock becomes obsolete. This explains why HSS may increase local robustness but will still perform worse under an obsolete-cost-only objective. Even in artificial perturbation scenarios, where demand is deliberately increased around the changeover moment, the no-HSS policy remains the best-performing solution. Larger perturbations can consume part of the available buffer and reduce obsolescence as seen in Section 6.4. However, the general pattern remains that HSS creates additional obsolete stock whenever the buffer is not fully used.

Another explanation for the limited performance is the aggregation of materials into material groups and the use of DOS as the HSS unit. Grouping materials by ABC class and demand pattern, in Section 4.3, makes the Sim-Opt procedure computationally feasible for an industry-scale MRP environment. At the same time, it also means that one HSS value is assigned to a heterogeneous set of items. Items within the same group can still differ substantially in MOQ, piece cost and material demand behaviour. Therefore, a group-level HSS policy may be useful for some items but harmful for others, with the aggregate objective masking such item-specific differences. Although, the perturbation size experiment in Section 6.3 also did not find improvements in objective value when looking at items individually. A limitation here is that only 5 items were tested in this experiment. Additionally, the material analysis showed that many changeover materials have intermittent or lumpy demand patterns. For these materials, a safety stock buffer can be especially valuable as it is harder to plan the procurement of these materials. But expressing the HSS in average daily demand is not always a strong representation of actual consumption behaviour. A buffer expressed in days of sales may therefore be too small to matter for some items, while already being excessive for others.

7.3. MRP-Opt structure and introducing hard safety stock

An important methodological contribution of this thesis is the introduction of the hard safety stock (HSS) parameter in Section 5.1. This intervention was necessary because the original safety stock constraint in the MRP-Opt model, described in Section 3.2.1, is a soft constraint. This means that stock may fall below the safety stock level by using a safety stock slack variable. The objective hierarchy prioritises obsolete stock minimisation above safety stock slack minimisation. Because of this priority, the original safety stock parameter could not serve as an external optimisation parameter. The proposed HSS parameter therefore represents a structural correction to make safety stock binding in the MRP-Opt model.

This design choice is defensible from a research perspective, but it should also be discussed in a critical way. By replacing a soft bound with a hard lower bound, this intervention introduces a stronger operational commitment to inventory coverage than the original MRP-Opt model was designed to respect. This may improve robustness, but it also reduces the flexibility of the MIP to resolve conflicting objectives in complex changeover situations. In environments where simultaneous or sequential changeovers already create combinatorial constraints, as described in Section 3.2, imposing a HSS constraint may force the model into more expensive decisions. At the same time, this intervention reveals something important about the currently used MRP tool itself. It shows that the manufacturer's existing MRP-Opt model, using a soft safety stock constraint, has only limited ability to value robustness under uncertainty. By allowing the stock levels to fall below the safety stock level, robustness is treated as an afterthought.

An additional observation concerns the different HSS assignment variants that were tested. The constant variant applies a hard safety stock throughout the entire planning horizon, the changeover weeks variant only activates the buffer during the item-specific changeover window, and the maximum changeover week variant restricts the intervention to the final permissible changeover week. Conceptually, the latter two variants were expected to provide a more targeted form of protection by limiting the amount of time a buffer is exposed to becoming obsolete. However, the experimental results showed only limited differences between the variants. Although the changeover-based variants consistently performed better than the constant variant, none of the assignment strategies outperformed the no-HSS baseline. This suggests that the timing of the HSS activation is less influential than the uncertainty for example. Because obsolete stock costs are prioritised and used as the sole performance metric in the Sim-Opt framework, any inventory retained through the HSS constraint is likely to be penalised once the changeover has occurred. The results therefore indicate that a more selective activation policy alone is insufficient to unlock the potential benefits of safety stock. The differences between the assignment variants do not demonstrate a superior activation strategy under the current modelling assumptions.

7.4. Uncertainty assignment approach

The approach for modelling the uncertainty presented in Chapter 4 and implemented in Section 5.7 is an important aspect of this thesis. The data analysis showed that nearly all observed variability in the production plans was demand-driven, while planning-induced variability was negligible. This justified the decision to model uncertainty primarily as stochastic changes in production volumes, rather than random reallocation decisions across locations and resources. The decomposition of these variances in Section 4.2.1 is therefore an important empirical validation of the uncertainty modelling in this study. At the same time, the limitations of this uncertainty representation should be recognised. The block-bootstrapping approach and the use of KDE-fitted zero-truncated distributions for the production plan deltas created a realistic image of the demand uncertainty. This mechanism was also used for introducing variability across the planning horizon in the Sim-Opt procedure. A hurdle model was also used for the uncertainty assignment as the original delta distributions were characterised by a very high share of zero deltas. The distributions also had long tails driven by a relatively small number of outliers. This means that the production plans were often stable, interrupted by occasional large changes in demand. This does raise the question of whether mean-based evaluation of obsolete stock costs across multiple simulations is sufficient to capture the relevance of rare disruptions in which safety stock might be most beneficial.

Another limitation concerns the time horizon of the available data. Section 4.1 explicitly notes that only six months of consistent weekly production plans were available. Within these six months, two weeks are also missing. The inadequate coverage of a full year could not reliably capture seasonal patterns or yearly campaign structures for example. This is a substantial limitation because in a FMCG environment, packaging changes and demand uncertainty are often strongly tied to annual events. While the uncertainty assignment approach developed in this thesis is appropriate for exploratory and methodological analysis, it may not be sufficient to support a definitive safety stock policy.

In addition to the realistic uncertainty assignment, this research also tested artificial uncertainty variants through deterministic perturbations in the production plan. These variants should not be interpreted as an attempt to represent the full stochastic behaviour observed in the historical data, but rather as controlled stress-test scenarios. By applying a fixed positive perturbation in either the changeover weeks or specifically in the maximum changeover week, the experiments isolate situations in which safety stock would theoretically be most valuable. This makes it possible to analyse whether HSS can protect against sudden demand increases around the most critical weeks of the changeover process. The results in Section 6.3 show that artificial perturbations can reduce obsolete stock in specific item-level cases when the additional demand is large enough to consume remaining material stock. However, they also confirm that the benefits of HSS are highly scenario-dependent. If the perturbation is too small, occurs outside the changeover window, or does not align with the material's remaining stock level, the imposed HSS remains unused and becomes obsolete. Therefore, the artificial uncertainty assignment is useful as a methodological stress test, but it should be interpreted separately from the realistic uncertainty approach. Its main value lies in showing under which controlled conditions safety stock could be more beneficial, rather than in providing a general representation of demand uncertainty.

7.5. Material grouping and material representation

The material classification used in this research provides a structured way to reduce the complexity of safety stock parametrisation. By combining the ABC analysis with the demand pattern classification in Section 4.3, the changeover materials are grouped according to both economic relevance and demand behaviour in Section 5.5. This is useful because safety stock is not expected to have the same relevance for every material. High-value materials have a larger impact on inventory costs, while materials with lumpy or intermittent demand are more difficult to plan because their demand is less regular. The resulting material groups therefore provide a practical basis for selecting candidate materials for the Sim-Opt framework. At the same time, this grouping is a simplification of reality. Items within the same material group can still differ in terms of MOQ, piece cost, production location and remaining demand before the changeover. Therefore, assigning one HSS value to an entire group can still mask item-level differences. A safety stock level that is suitable for one material in a group may be too high or too low for another material in the same group.

A limitation of the experiments is that they were only performed on materials from one material group, MG2. This group was selected because A-class materials with lumpy demand were expected to be among the most promising candidates for safety stock optimisation: they combine high economic relevance with uncertain demand quantities. However, testing only one group limits the generalisability of the results. The finding that HSS does not improve obsolete stock costs for MG2 does not necessarily mean that the same conclusion holds for all other material groups. For example, intermittent materials or smooth high-value materials may respond differently to HSS. The perturbation size and proof-of-concept experiments focused on an even smaller subset of materials in MG2, making their results even more specific to individual items. This is an important limitations, however the trend seen in their results is similar to the variants experiments, which were tested on more materials. The addition of safety stock always increases the obsolete material costs. This suggests that the reason behind these findings is not primarily caused by the material selection, but rather by the current objective function and Sim-Opt setup.

7.6. Optimisation method

The use of Full Enumeration technique, described in Section 5.4, as the optimisation method provided a clear and reliable way to evaluate the tested HSS parameter range, because every integer value between the lower and upper bound was evaluated. This made the results easy to interpret and ensured that, within the defined search space of 0 to 7 days of sales, the best-performing HSS value was found. However, this technique also created several limitations. Full Enumeration is only feasible when the number of parameters and tested values remains small, because the number of required Sim-Opt evaluations increases rapidly with wider HSS bounds or when more optimisation parameters are added. As a result, the experiments had to remain relatively limited in scope, and the method was mainly suitable for testing one material group and a small set of HSS values. The results show that the lowest HSS values consistently performed best, which suggests that expanding the search range above 7 DOS would likely not have improved the objective value under the current obsolete cost objective. Nevertheless, Full Enumeration limits the scalability of the framework and is not suitable for larger, more realistic parameter spaces.

7.7. Contribution to literature

In relation to the literature reviewed in Chapter 2, this thesis contributes in several ways. First, it explicitly focuses on changeover materials, whereas much of the existing literature on safety stock parametrisation under uncertainty concentrates on stable materials. Secondly, this study treats the manufacturer's MRP process as a sequential decision analytics problem. It evaluates parameter choices over multiple weekly re-optimisations rather than a single deterministic run. Third, it applies a simulation-based optimisation framework to an existing large-scale industrial problem, while treating the model as a black-box. This combination, which is identified in Section 2.3, is relatively uncommon in literature.

The academic value of the thesis does not depend solely on demonstrating a positive cost improvement. This study also contributes by identifying the conditions under which safety stock parametrisation is difficult to make effective. Examples of this are partial objective functions and strong internal model preferences for lean inventories. Building a working simulation-based optimisation framework through which further research on safety stock optimisation could be performed, also provides academic value. The limitations discussed in this chapter present relevant insights for both researchers and practitioners, and make the thesis valuable even though the final business result was not positive.

7.8. Practical implications for the manufacturer

For the manufacturer, the main implication of this research is that the developed framework should not yet be regarded as a final decision rule for the setting of safety stock levels of changeover materials. Although the study demonstrates that safety stock parameters can be optimised around the existing MRP-Opt model in a stochastic simulation environment, the final results do not show a performance improvement caused by HSS under the KPI currently used in this research. A generic rollout of hard safety stock values for all changeover materials is therefore not recommended on the basis of the present findings alone. Instead, the framework should currently be interpreted as a structured decision-support instrument that can support targeted experimentation and learning.

At the same time, the results indicate that the current planning logic can undervalue robustness. The original MRP-Opt structure prioritises obsolete stock minimisation and allows safety stock violations through a soft constraint. Due to this, inventories of changeover materials can still be planned below their intended safety stock level when this benefits the obsolete stock sub-objective. This suggests that the economic consequences of insufficient robustness are not yet fully visible in the current MRP-Opt objective structure. The manufacturer is therefore recommended to improve the measurement and availability of data such as inventory holding costs and costs associated with rush orders. With this information available, future analyses can evaluate the full economic trade-off of robustness rather than doing performance measurements with obsolete materials only.

Beyond the safety stock optimisation results, this thesis creates value for the manufacturer through the data analysis that was conducted on both materials and production plans. The production plan analysis demonstrated that more than 99.98% of the observed variability originates from demand fluctuations rather than planning decisions. This suggests that future improvement initiatives could focus on managing demand uncertainty instead of adjusting production allocation logic. Furthermore, the developed methodology for analysing historical production plans using block bootstrapping and kernel density estimation provides the manufacturer with a structured and reproducible way to quantify uncertainty and evaluate the robustness of planning policies. The material analysis also delivers direct practical insights. By classifying changeover materials according to both demand behaviour (smooth, intermittent, erratic and lumpy) and economic importance through an ABC analysis, the manufacturer gains a more detailed understanding of its material and inventory portfolio. This classification highlights which materials contribute most to inventory value and which materials are characterised by high uncertainty. This creates a valuable decision-support foundation that can be reused for future planning initiatives.

This thesis provides the manufacturer with a valuable methodological asset: a simulation-based optimisation framework that can be reused to evaluate other planning parameters or policy designs in the future. Although the specific safety stock policy tested in this thesis did not outperform the baseline, the Sim-Opt framework is still useful. It can be used to test other robustness interventions in a rolling-horizon setting without analytically reformulating the underlying MIP model. A practical next step would be to apply the framework to other scenarios. The findings of this study suggest that the benefits of safety stock are likely to be material-specific rather than system-wide. Some materials may benefit from safety stocks to mitigate uncertainty, while others simply experience additional obsolescence risk caused by these buffers. The manufacturer is therefore advised to use the current framework to identify candidate materials or material groups for which HSS can work in the future. The simulation-based optimisation framework can support such future analyses. At this stage, it should primarily be regarded as a methodological foundation for structured experimentation, rather than as a stand-alone operational policy.

8 Conclusion

This chapter first presents the conclusions of this research, by answering the research questions in Section 1.3. After presenting the conclusions, several recommendations for future research will be elaborated on in this chapter.

8.1. Conclusion

This thesis set out to examine how simulation-based optimisation can be used to parametrise safety stocks of changeover materials in a deterministic MIP model, such that rolling horizon performance improves in a stochastic MRP environment. Based on the performed research, the following main conclusion is established. This research proved that simulation-based optimisation can indeed be used to determine safety stock levels for changeover materials in such an environment, but this did not lead to an improvement in rolling horizon performance within the scope of this study. In that sense, the developed framework proved to be methodologically feasible, but its performance fell short compared to the no HSS baseline.

Regarding RQ1, this thesis finds that the state of the art on determining planning parameters in MRP under uncertainty is mainly focused on demand uncertainty and on standard materials. The reviewed literature shows that uncertainty in MRP is typically mitigated through buffers such as safety stock and safety lead time. More recent studies increasingly use simulation-based optimisation, metaheuristics or stochastic optimisation methods to determine planning parameters. Only limited attention is paid to changeover materials specifically. This is relevant, since changeover materials are characterised by a different trade-off than standard materials. Holding additional stock may improve robustness against uncertainty in the production plan, but it also increases the risk of material obsolescence when a material is phased out. This thesis therefore confirms that a research gap exists with regard to the parametrisation of safety stocks for changeover materials in stochastic MRP environments. Meanwhile, the application on an industry-size problem provides an additional contribution to the literature.

RQ2 concerned the current state of the manufacturer's MRP process. The analysis showed that this process can be identified as a sequential decision analytics process in which the MRP-Opt model is solved every week using updated data. The current safety stock logic in the manufacturer's MRP process is deterministic and the corresponding safety stock constraint in the MRP-Opt model is not bounding. This allows stock levels to fall below the specified safety stock level whenever that helps reduce obsolete stock costs. As obsolete stock minimisation has a higher priority in the model, robustness is effectively treated as secondary. Because of this, changeover materials are often planned with limited protection against changes in demand. The data analysis showed that almost all observed variability in the production plans is demand-driven, while location planning-induced variability is negligible. The material analysis further demonstrated that the manufacturer's changeover materials are dominated by intermittent and lumpy demand patterns. In addition, the ABC analysis showed a strong concentration of inventory value within a relatively small subset of items, with class A materials accounting for the majority of the total value. These findings motivated the material grouping approach used in the Sim-Opt framework and highlighted A-class lumpy and intermittent materials as the most relevant candidates for safety stock optimisation.

To answer RQ3, a simulation-based optimisation framework was developed in which the MRP-Opt model was treated as a black box. The MRP-Opt tool was embedded in the Sim-Opt framework, which used an Optimisation with Simulation-based Iterations approach. In this framework, hard safety stock values were introduced as binding parameters. Meanwhile, uncertainty was assigned to the production plans using historical production-plan deltas derived through a method based on block bootstrapping. The safety stock parameter space was searched using a Full Enumeration technique. Meanwhile, the framework explored several variants of applying HSS and also more artificial uncertainty variants. The experimental results showed that simulation-based optimisation can be used to optimise safety stock parameters in a rolling horizon setting, while imitating the manufacturer's MRP process using its MRP-Opt model as it is used in practice.

The experimental results show that the proposed hard safety stock approach did not produce lower obsolete stock costs than the baseline situation across all tested scenarios. In the realistic uncertainty experiments, as well as in the artificial perturbation scenarios designed to create favourable conditions for safety stock, low safety stock levels consistently generated the best performance. Increasing the hard safety stock generally resulted in higher obsolete stock costs, while the Sim-Opt optimisation procedure repeatedly converged towards to the lowest buffer levels. The experiments further demonstrated that the framework produced stable and reproducible results under different Sim-Opt settings, uncertainty assignments and HSS magnitudes and application time windows. None of the experiments was able to change the finding that additional safety stock was unable to compensate for uncertainty. These findings indicate that, within the tested environment and performance measure, retaining extra inventory around changeovers introduces more obsolete material than it prevents. This makes a generic hard safety stock policy ineffective for improving system-wide performance. The conclusion of this thesis is therefore not that safety stock is irrelevant for changeover materials. It is rather that no generalisable safety stock policy could be found to provide an improvement under the current modelling assumptions, limitations and available data.

In conclusion, this thesis shows that simulation-based optimisation is a viable method for parametrising safety stocks of changeover materials in a deterministic MIP-based MRP environment. Such an approach can be implemented in a realistic industrial setting while treating the original MRP optimisation model as a black box. However, the results show that this technical capability does not automatically lead to improved performance in the current research setup. The main contribution of this thesis therefore lies in the developed framework itself. Other contributions lie in the insights from the data analyses and in making explicit where current KPI design, data availability and model design limit the measurable value of robustness. The Sim-Opt framework developed in this study should be seen as a foundation for further research and future practical testing, rather than as a final decision rule for safety stock policy in the manufacturer's MRP process.

8.2. Recommendations for further research

This research showed that simulation-based optimisation can be used to parametrise hard safety stock levels for changeover materials in a deterministic, black-box MRP optimisation model. However, the experiments did not show an improvement in rolling-horizon performance when performance was measured only through obsolete stock costs. Across the tested uncertainty scenarios, higher hard safety stock levels generally increased obsolete material costs, while the lowest tested safety stock levels consistently performed best. This suggests that a generic hard safety stock policy is not effective under the current objective structure. Future research should therefore focus on improving how robustness, uncertainty and material-specific trade-offs are represented in the framework. The following sections will explore specific recommendations for further research.

8.2.1. Broaden the objective function

The most important recommendation is to extend the objective function beyond obsolete stock costs. In this thesis, obsolete stock costs were used as the only optimisation objective because inventory holding costs and rush-order costs were not available from MRP-Opt. However, this only captures the downside of safety stock. The potential benefits, such as preventing last-minute procurement, avoiding MOQ-triggered orders or reducing operational disruption, are not measured. As a result, the optimisation is naturally pushed toward low or zero safety stock levels.

Future research should include obsolete stock costs, inventory holding costs and rush-order costs in one objective function. For example, a policy with one or two days of sales as safety stock may create slightly more obsolete stock after a changeover, but it could still be beneficial if it prevents an expensive last-minute MOQ order shortly before the changeover. This could lead to different conclusions than in this research, where safety stock could only be penalised through additional obsolete material.

To support a broader objective function, more operational data should be collected, analysed, and potentially integrated. Relevant data includes historical rush orders, urgent transport costs, inventory holding costs and write-off costs. These data sources would make it possible to evaluate whether safety stock prevents costs that are currently invisible in the simulation output. This is especially relevant because the current results should not necessarily be interpreted as proof that safety stock has no value. Rather, they show that safety stock does not reduce obsolete stock costs when evaluated in isolation. With richer cost data, future research could distinguish between safety stock that is truly inefficient and safety stock that is useful.

8.2.2. Move toward material-specific safety stock policies

The material analysis showed that changeover materials differ strongly in both economic importance and demand behaviour. Many items have intermittent or lumpy demand, while the ABC analysis showed that a relatively small share of items represents most of the total inventory value. This means that a single generic safety stock policy is unlikely to work well for all materials.

Future research should therefore test more material-specific or scenario-specific policies. A promising direction is to focus on high-value materials with uncertain demand, such as A-lumpy and A-intermittent materials. These materials have both financial relevance and demand uncertainty, making them more likely to benefit from targeted robustness measures. Although these kinds of materials were tested in this research, it did test all of them. Meanwhile, a very limited number of materials was tested individually to see what the effect of HSS and uncertainty had on the simulation results. Therefore, it is recommended to test more items individually. At the same time, testing items in other material groups could show if HSS is beneficial for items with different material characteristics. For example, an A-lumpy material with a high MOQ and a short changeover window may justify a small buffer, while a low-value intermittent material may not.

8.2.3. Refine the uncertainty modelling

The uncertainty analysis showed that almost all observed variability in the production plans was demand-driven rather than caused by planning reallocations. This supports the decision to model uncertainty through demand perturbations. However, the available historical data covered only about half a year, which limits the ability to capture seasonality, campaign effects and rare demand shocks. Future research should repeat the uncertainty analysis with at least one full year, and preferably multiple years, of production plan data. This would make the fitted uncertainty distributions more reliable. Another recommendation would be to perform an extra analysis into the distinction between normal demand variation, seasonal patterns and exceptional disruptions. This distinction creates the possibility for the HSS policy to anticipate on seasonal demand patterns, and to disregard extreme disruptions. With this, the HSS parametrisation could become more accurate and fit better with realistic demand patterns.

Because changeover materials become obsolete after transition, the timing of uncertainty is especially important. A demand increase shortly before the final changeover week can create a very different risk than a demand increase much earlier in the horizon. The artificial perturbation experiments already explored this idea, but future research could model uncertainty around changeover windows in a more data-driven way. For example, historical changeovers could be analysed to determine whether certain materials or product groups often experience late demand increases before changeover. Safety stock could then be activated only for those cases, instead of applying it across broader time windows or material groups. This may reveal situations in which a small, temporary buffer is useful, even though the broader hard safety stock policies tested in this thesis did not improve performance.

8.2.4. Explore robustness measures beyond safety stock

The results indicate that hard safety stock alone may be too limited as a robustness measure for changeover materials. Since additional stock can quickly become obsolete after a changeover, future research should compare safety stock with other interventions, such as safety lead times, adjusted order timing or MOQ-aware replenishment rules. For instance, instead of holding extra material, the system could delay procurement until more reliable demand information is available or split orders when suppliers allow it. These alternatives may provide robustness without creating as much leftover inventory. A criterion for this is that the necessary data should be available to explore these options. Meanwhile, some of these alternatives might also require too many fundamental changes to the manufacturer's MRP process.

8.2.5. Improve computational scalability

Full Enumeration was useful in this thesis because the tested parameter space was small and it provided a clear benchmark. However, it is not suitable for larger-scale optimisation with more materials, wider parameter ranges or multiple decision variables. Future research could therefore explore more scalable optimisation methods, such as Simulated Annealing. Simulated Annealing would create the possibility of optimising multiple decision variables at a time, while also exploring a larger parameter space. When other changes to the Sim-Opt framework provide scenarios in which HSS is found to be beneficial, it is worth exploring the HSS parameters on a more granular level, for example 2 decimals, compared to integers. This increases the parameter space, requiring a more scalable optimisation method than Full Enumeration. Improving scalability would allow the framework to evaluate larger and more realistic material sets. It would also make it possible to optimise multiple material groups or planning parameters simultaneously, rather than testing one limited safety stock setting at a time.

8.2.6. Validate the framework with practical planning cases

Finally, future research should validate the framework using historical planning cases and planner feedback. Planners often know when the MRP output is theoretically cost-efficient but operationally risky. Comparing the framework's recommendations with real planner interventions, rush-order cases or obsolete stock incidents would help assess whether the model captures the right trade-offs. This would also help determine whether the framework can become a practical decision-support tool. Even if a policy does not minimise obsolete stock costs, it may still be valuable if it reduces operational risk or prevents difficult last-minute planning decisions.

Overall, future research should build on the framework developed in this thesis, but should not directly adopt the tested hard safety stock policies as final planning rules. The framework is most valuable as a structured method for evaluating planning parameters under uncertainty.

References

1 Grinvalds, K. (n.d.). *Generic warehouse interior with high shelves and boxes of products for distribution* [Photograph. Adobe Stock. Accessed 14 September 2026]. https://stock.adobe.com/nl/images/generic-warehouse-interior-with-high-shelves-and-boxes-of-products-for-distribution/481003062?continue-checkout=1&isa0=1&state=%7B%22ac%22%3A%22stock.adobe.com%22%7D

2 Dolgui, A., & Prodhon, C. (2007). Supply planning under uncertainties in MRP environments: A state of the art. *Annual Reviews in Control*, *31*(2), 269–279. https://doi.org/10.1016/j.arcontrol.2007.02.007

3 Javadi, S. M., Sadjadi, S. J., Teimoury, E., & Makui, A. (2025). Material requirements planning with a novel lot sizing method and a new algorithm for production scheduling. *Scientific Reports*, *15*(1). https://doi.org/10.1038/s41598-025-13197-8

4 Bill of Materials (BOM) Guide | Prometheus Group [Retrieved on 4th of March 2026]. (2025, November). https://www.prometheusgroup.com/learning-center/what-is-a-bill-of-materials-bom

5 Lim, W. S., & Tang, C. S. (2006). Optimal product rollover strategies. *European Journal of Operational Research*, *174*(2), 905–922. https://doi.org/10.1016/j.ejor.2005.04.031

6 Liang, C., Çakanyildirim, M., & Sethi, S. P. (2014). Analysis of product rollover strategies in the presence of strategic customers. *Management Science*, *60*(4), 1033–1056. https://doi.org/10.1287/mnsc.2013.1803

7 King, P. L. (2011). The CoDe CraCK: Understanding safety stock and mastering its equations. *APICS magazine*, 33–35. https://web.mit.edu/2.810/www/files/readings/King_SafetyStock.pdf

8 Powell, W. B. (2019, June). A unified framework for stochastic optimization. https://doi.org/10.1016/j.ejor.2018.07.014

9 Figueira, G., & Almada-Lobo, B. (n.d.). *Hybrid Simulation-Optimization Methods: A Taxonomy and Discussion* (tech. rep.).

10 Jodlbauer, H., & Huber, A. (2008). Service-level performance of MRP, kanban, CONWIP and DBR due to parameter stability and environmental robustness. *International Journal of Production Research*, *46*(8), 2179–2195. https://doi.org/10.1080/00207540600609297

11 Gungor, Z. E., Inan, G. G., Marques-McEwan, M., Evans, S., & Bititci, U. S. (2025). Reimagining changeovers as a key value-adding business capability. *Production Planning and Control*. https://doi.org/10.1080/09537287.2025.2554163

12 Fiorotto, D. J., Huaccha Neyra, J. d. C., & de Araujo, S. A. (2020). Impact analysis of setup carryover and crossover on lot sizing problems. *International Journal of Production Research*, *58*(20), 6350–6369. https://doi.org/10.1080/00207543.2019.1680892

13 Tordecilla, R. D., Juan, A. A., Montoya-Torres, J. R., Quintero-Araujo, C. L., & Panadero, J. (2021, January). Simulation-optimization methods for designing and assessing resilient supply chain networks under uncertainty scenarios: A review. https://doi.org/10.1016/j.simpat.2020.102166

14 Angkiriwang, R., Pujawan, I. N., & Santosa, B. (2014). Managing uncertainty through supply chain flexibility: reactive vs. proactive approaches. *Production and Manufacturing Research*, *2*(1), 50–70. https://doi.org/10.1080/21693277.2014.882804

15 Barros, J., Cortez, P., & Carvalho, M. S. (2021, January). A systematic literature review about dimensioning safety stock under uncertainties and risks in the procurement process. https://doi.org/10.1016/j.orp.2021.100192

16 Hong, Z., Lee, C. K., & Zhang, L. (2018, September). Procurement risk management under uncertainty: a review. https://doi.org/10.1108/IMDS-10-2017-0469

17 Seiringer, W., Altendorfer, K., Felberbauer, T., Bokor, B., & Brockmann, F. (2025). How periodic forecast updates influence MRP planning parameters: A simulation study. *Simulation Modelling Practice and Theory*, *142*. https://doi.org/10.1016/j.simpat.2025.103115

18 Werth, B., Karder, J., Beham, A., & Altendorfer, K. (2022). Simulation-based Optimization of Material Requirements Planning Parameters. *Procedia Computer Science*, *217*, 1117–1126. https://doi.org/10.1016/j.procs.2022.12.310

19 Seiringer, W., Castaneda, J., Altendorfer, K., Panadero, J., & Juan, A. A. (2022). Applying Simheuristics to Minimize Overall Costs of an MRP Planned Production System. *Algorithms*, *15*(2). https://doi.org/10.3390/a15020040

20 Seiringer, W., Altendorfer, K., Castaneda, J., Panadero, J., & Juan, A. A. (2021). Applying Simheuristics for Safety Stock and Planned Lead Time Optimization in a Rolling Horizon MRP System under Uncertainty. *Proceedings - Winter Simulation Conference*, *2021-December*. https://doi.org/10.1109/WSC52266.2021.9715294

21 Thevenin, S., Adulyasak, Y., & Cordeau, J. F. (2021). Material Requirements Planning Under Demand Uncertainty Using Stochastic Optimization. *Production and Operations Management*, *30*(2), 475–493. https://doi.org/10.1111/poms.13277

22 Gansterer, M., Almeder, C., & Hartl, R. F. (2014). Simulation-based optimization methods for setting production planning parameters. *International Journal of Production Economics*, *151*, 206–213. https://doi.org/10.1016/j.ijpe.2013.10.016

23 Graves, S. C., & Willems, S. P. (2008). Strategic inventory placement in supply chains: Nonstationary demand. *Manufacturing and Service Operations Management*, *10*(2), 278–287. https://doi.org/10.1287/msom.1070.0175

24 Campuzano-Bolarín, F., Mula, J., Díaz-Madroñero, M., & Legaz-Aparicio, Á. G. (2020). A rolling horizon simulation approach for managing demand with lead time variability. *International Journal of Production Research*, *58*(12), 3800–3820. https://doi.org/10.1080/00207543.2019.1634849

25 Badakhshan, E., Mustafee, N., & Bahadori, R. (2024). Application of simulation and machine learning in supply chain management: A synthesis of the literature using the Sim-ML literature classification framework. *Computers and Industrial Engineering*, *198*. https://doi.org/10.1016/j.cie.2024.110649

26 Deliu, N. (2024). Reinforcement learning for sequential decision making in population research. *Quality and Quantity*, *58*(6), 5057–5080. https://doi.org/10.1007/s11135-023-01755-z

27 Wang, H., Tao, J., Peng, T., Brintrup, A., Kosasih, E. E., Lu, Y., Tang, R., & Hu, L. (2022). Dynamic inventory replenishment strategy for aerospace manufacturing supply chain: combining reinforcement learning and multi-agent simulation. *International Journal of Production Research*, *60*(13), 4117–4136. https://doi.org/10.1080/00207543.2021.2020927

28 Jalali, H., & Van Nieuwenhuyse, I. (n.d.). *Simulation optimization in inventory replenishment: A classification* (tech. rep.).

29 Fan, W., Hong, L. J., Jiang, G., & Luo, J. (2024). Review of Large-Scale Simulation Optimization. http://arxiv.org/abs/2403.15669

30 Schwartz, J. D., Wang, W., & Rivera, D. E. (2006). Simulation-based optimization of process control policies for inventory management in supply chains. *Automatica*, *42*(8), 1311–1320. https://doi.org/10.1016/j.automatica.2006.03.019

31 Rojas-Gonzalez, S., & Van Nieuwenhuyse, I. (2020, April). A survey on kriging-based infill algorithms for multiobjective simulation optimization. https://doi.org/10.1016/j.cor.2019.104869

32 Jalali, H., Van Nieuwenhuyse, I., & Picheny, V. (2017). Comparison of Kriging-based algorithms for simulation optimization with heterogeneous noise. *European Journal of Operational Research*, *261*(1), 279–301. https://doi.org/10.1016/j.ejor.2017.01.035

33 Ghafariasl, P., Mahmoudan, A., Mohammadi, M., Nazarparvar, A., Hoseinzadeh, S., Fathali, M., Chang, S., Zeinalnezhad, M., & Garcia, D. A. (2024). Neural network-based surrogate modeling and optimization of a multigeneration system. *Applied Energy*, *364*. https://doi.org/10.1016/j.apenergy.2024.123130

34 Katoch, S., Chauhan, S. S., & Kumar, V. (2021). A review on genetic algorithm: past, present, and future. *Multimedia Tools and Applications*, *80*(5), 8091–8126. https://doi.org/10.1007/s11042-020-10139-6

35 Fontes, D. B., Homayouni, S. M., & Gonçalves, J. F. (2023). A hybrid particle swarm optimization and simulated annealing algorithm for the job shop scheduling problem with transport resources. *European Journal of Operational Research*, *306*(3), 1140–1157. https://doi.org/10.1016/j.ejor.2022.09.006

36 Zabraoui, O., Hmamou, Y., Chafi, A., & Alami, S. K. (2025). A comparative study of multi-algorithm optimization for inventory analytics in supply chains. *Supply Chain Analytics*, *12*. https://doi.org/10.1016/j.sca.2025.100154

37 Scala, P., Mota, M. M., Wu, C. L., & Delahaye, D. (2021). An optimization–simulation closed-loop feedback framework for modeling the airport capacity management problem under uncertainty. *Transportation Research Part C: Emerging Technologies*, *124*. https://doi.org/10.1016/j.trc.2020.102937

38 Maitra, S. (n.d.). *A System-Dynamic Based Simulation and Bayesian Optimization for Inventory Management* (tech. rep.).

39 Fachrizal, R., Shepero, M., Van Der Meer, D., Munkhammar, J., & Widén, J. (2020). Smart charging of electric vehicles considering photovoltaic power production and electricity consumption: A review. *eTransportation*, *4*, 100056. https://doi.org/10.1016/j.etran.2020.100056

40 Mader, M., Mader, W., Sommerlade, L., Timmer, J., & Schelter, B. (2013). Block-bootstrapping for noisy data. *Journal of Neuroscience Methods*, *219*(2), 285–291. https://doi.org/10.1016/j.jneumeth.2013.07.022

41 Torres, M. I., Lupton, R., Marsh, E., Srubar, W. V., & Allen, S. (2026). Using kernel density estimation and the Dirichlet distribution for uncertainty quantification of building material emissions. *Resources, Conservation and Recycling*, *234*. https://doi.org/10.1016/j.resconrec.2026.109022

42 Zeileis, A., Kleiber, C., & Jackman, S. (n.d.). *Regression Models for Count Data in R* (tech. rep.). http://CRAN.R-project.org/package=pscl.

43 Syntetos, A. A., Boylan, J. E., & Croston, J. D. (2005). On the categorization of demand patterns. *Journal of the Operational Research Society*, *56*(5), 495–503. https://doi.org/10.1057/palgrave.jors.2601841

44 Qaffas, A. A., Hajkacem, M.-A. B., Ben Ncir, C.-E., & Nasraoui, O. (n.d.). Interpretable Multi-criteria ABC analysis based on semi-supervised clustering and Explainable Artificial Intelligence. https://doi.org/10.1109/ACCESS.2022.0092316

45 Pham, E. (2026). Safety Stock Calculation: Mathematical Approach to Demand Uncertainty. [Retrieved on 8th of June 2026]. https://resources.rework.com/libraries/manufacturing-growth/safety-stock-calculation

A
AI statement

During the preparation of this thesis, Artificial Intelligence (AI) tools were used as supportive aids in several parts of the research and writing process. AI was used to assist with the development and refinement of Python code for modelling and experimentation, to support parts of the data analysis, and to help structure and report the results. In addition, AI was used to improve the readability of the text and strengthen the use of scientific and formal academic language. AI was used only as a support tool and not to independently provide results. Any AI-generated suggestions, code fragments or text revisions were critically reviewed, verified and adjusted before being included.

B
MRP-Opt inputs & outputs

B.1. MRP-Opt input

The model uses a large amount of data to optimise the changeover process. It uses three main categories of input data: i) Master data; ii) Periodic data and iii) Vendor data. The optimisation tool is connected to the manufacturer's API, so it automatically uploads the available data on a daily or weekly basis. Some specific data, for example changeover data, needs to be added manually by the manufacturer's planners. The manufacturer relies on vendors to supply or update their data in a timely manner. Table B.1 gives an overview of all the input data that the model requires.

Table B.1: Overview of data inputs and their descriptions

Data Input	Description
location master	location description, local and global location IDs, location type, and country
item master	item description, global item ID, and common material group (CMG) code
local item mapping	mapping between local item codes and global item codes
conversion factors	volume-based conversion factors for finished products
location stock	stock levels per location
scheduled receipts	process orders or scheduled receipts for finished product production
BOM	bill of materials for finished products
production step	mapping between finished product–resource combinations and production methods
production method	mapping of production methods to BOM versions
purchase order	purchase orders and call-offs
vendor quotas	vendor quota splits for materials
material plant	material–plant parameters such as safety stock, minimum lot size, and piece costs
goods receipt	received purchase orders
goods movements	goods movement records for items
vendor master	mapping of vendor codes to vendor names
production plan	production plan for finished products

B.2. MRP-Opt output

The tool will send an optimisation request to a server, after which it returns the optimisation results after several hours. The results are visualised in an informative user interface, so planners are able to work with it easily. Table B.2 mentions the output that the optimisation generates and what they mean. For example, the model generates a changeover projection, from which planners can see what the projected stock levels of changeover materials will be over the coming months.

Table B.2: Overview of optimisation output and their descriptions

Model output	Description
vendor stock projection	overview per vendor of all items and their stock projections for all relevant items
national stock projection	projection of stock and forecast at a national level, aggregated over all vendors and production plants
changeover projection	projection of stock levels for items that have planned changeovers
obsolete costs	overview of obsolete costs resulting from planned changeovers, including obsolete quantities and associated costs for old materials
obsolete analysis	detailed analysis of items becoming obsolete due to planning decisions
location stock projection	aggregated stock projection at the location level for each material
stock coverage	overview of stock coverage for relevant materials
deadstock	items marked as obsolete can be flagged as deadstock, making them part of obsolete cost calculations and allowing the model to disregard safety stock to minimize obsoletes
material restriction	overview of material restrictions impacting supply, production, or planning
optimisation summary	summary of the optimisation results across the supply chain
vendor forecast	forecast information for suppliers, including purchase orders, call-offs, and model-generated forecast

C
Statistical metrics

C.1. Sample standard deviation

$$s = \sqrt{\frac{1}{n-1}\sum_{i=1}^{n}(X_i - \bar{X})^2}$$

Where:

- s = sample standard deviation

- X_i = i-th observation in the dataset
- $\bar{X}$ = sample mean
- n = number of observations
- $\sum$ = summation over all observations

C.2. Standard error

$$\text{SE} = \frac{s}{\sqrt{n}}$$

Where:

- SE = standard error

C.3. Coefficient of variation

$$\mathrm{CV} = \frac{s}{\bar{X}} \times 100\%$$

Where:

- CV = coefficient of variation

C.4. Delta

$$\Delta x = x_{\text{current}} - x_{\text{previous}}$$

Where:

- Δx = change in variable x
- x_{current} = current value of x
- x_{previous} = previous value of x

D
Delta distributions per block

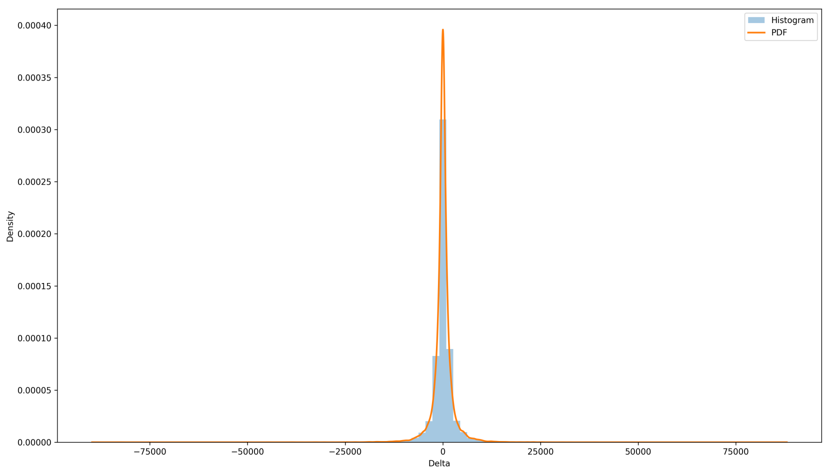

Figure D.1: Zero-truncated delta distribution with PDF of **block 1**

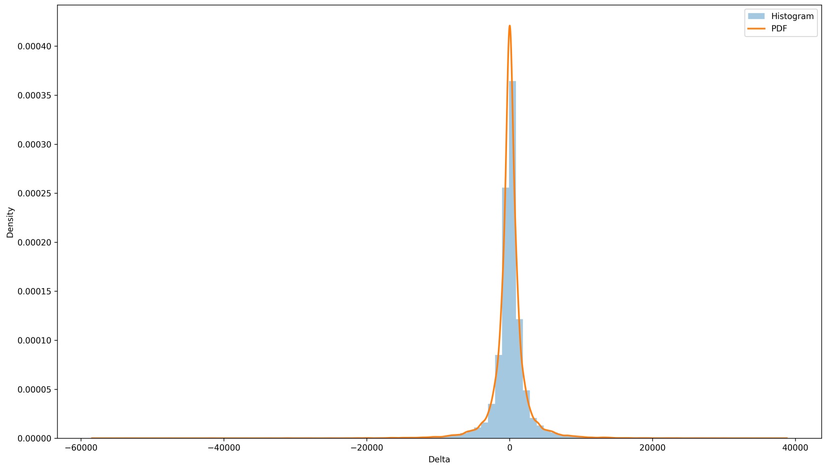

Figure D.2: Zero-truncated delta distribution with PDF of **block 2**

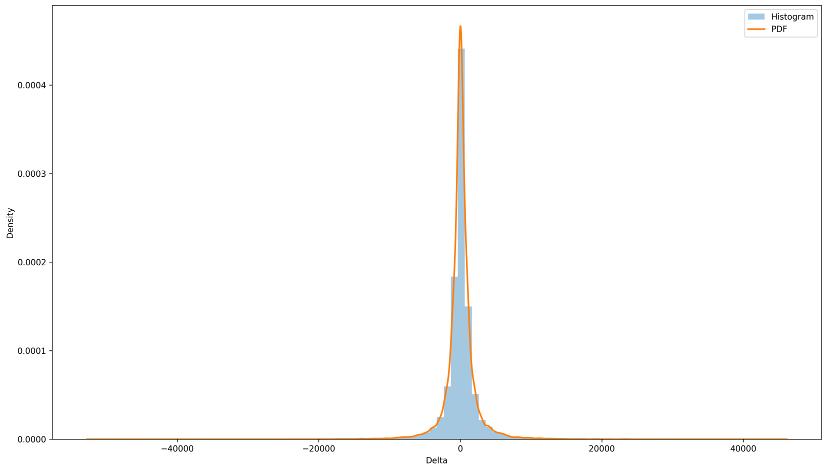

Figure D.3: Zero-truncated delta distribution with PDF of **block 3**

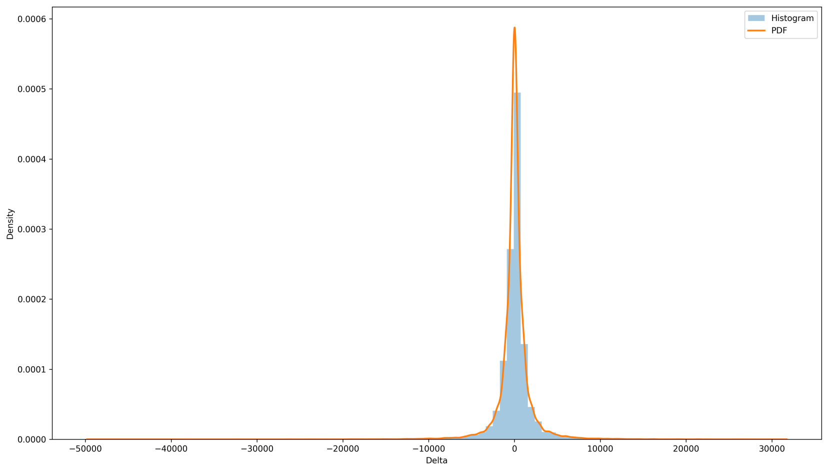

Figure D.4: Zero-truncated delta distribution with PDF of **block 4**

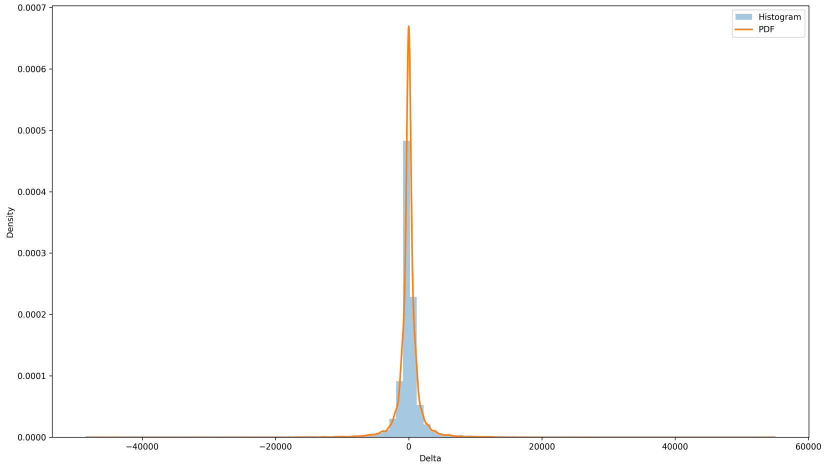

Figure D.5: Zero-truncated delta distribution with PDF of **block 5**

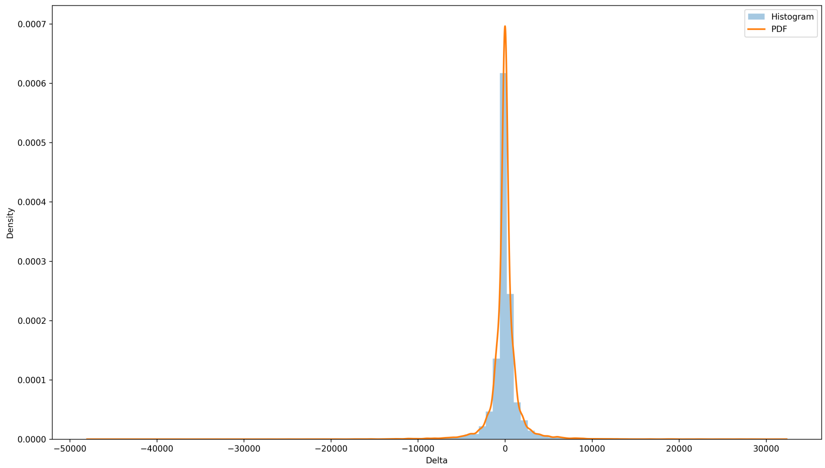

Figure D.6: Zero-truncated delta distribution with PDF of **block 6**

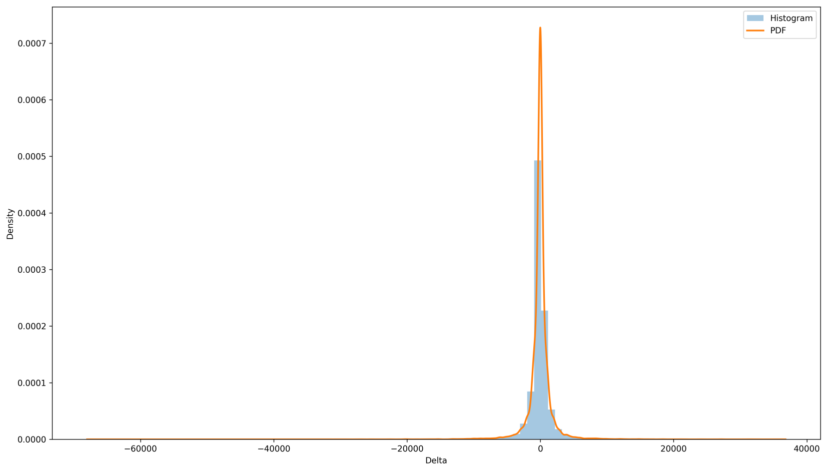

Figure D.7: Zero-truncated delta distribution with PDF of **block 7**

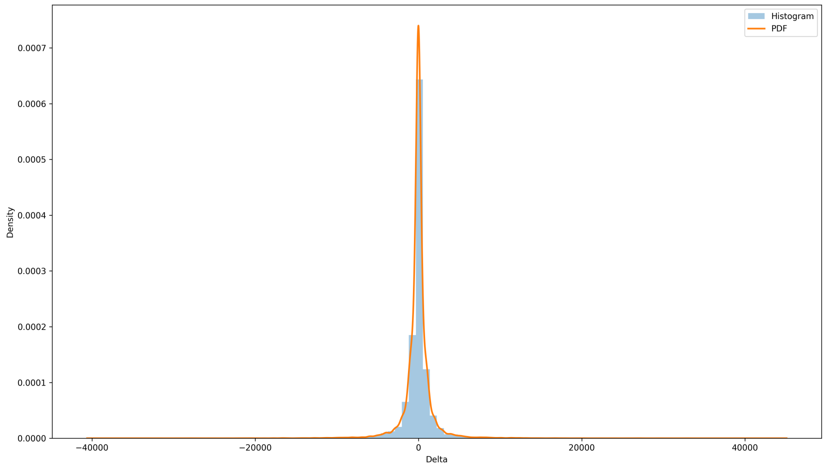

Figure D.8: Zero-truncated delta distribution with PDF of **block 8**

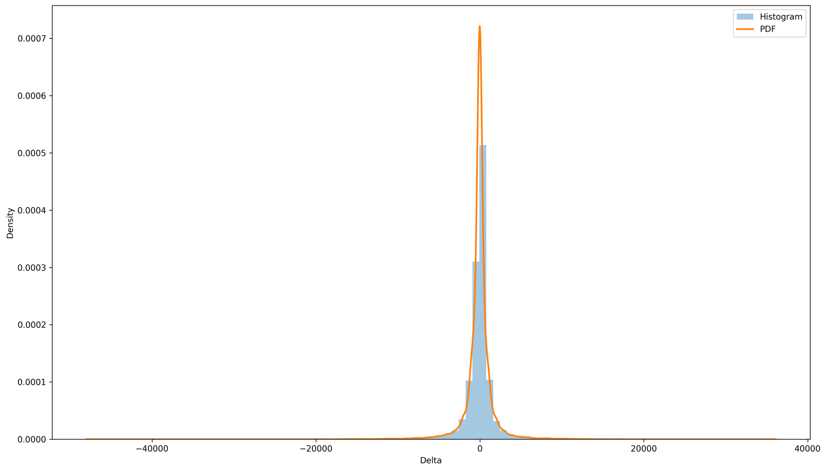

Figure D.9: Zero-truncated delta distribution with PDF of **block 9**

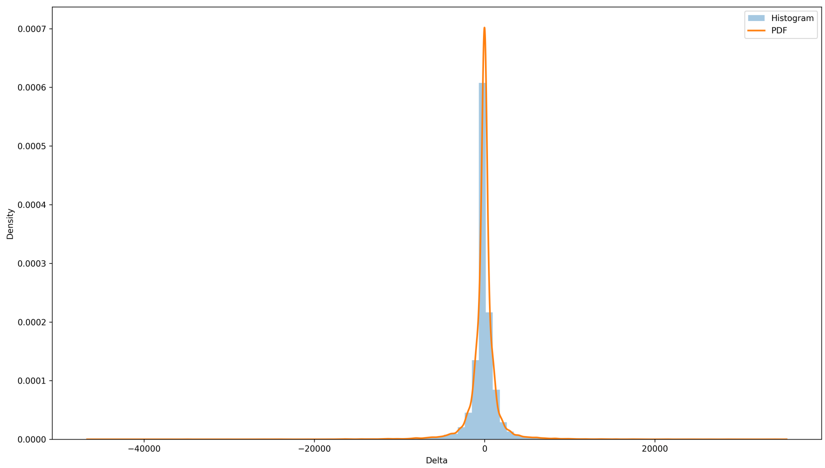

Figure D.10: Zero-truncated delta distribution with PDF of **block 10**

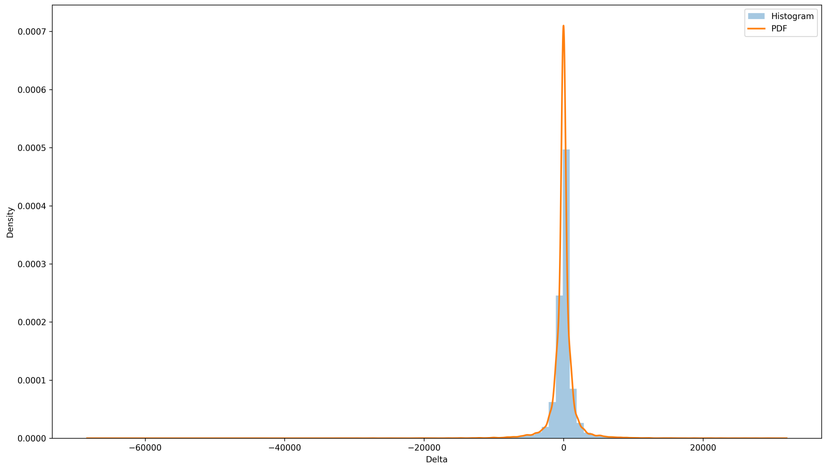

Figure D.11: Zero-truncated delta distribution with PDF of **block 11**

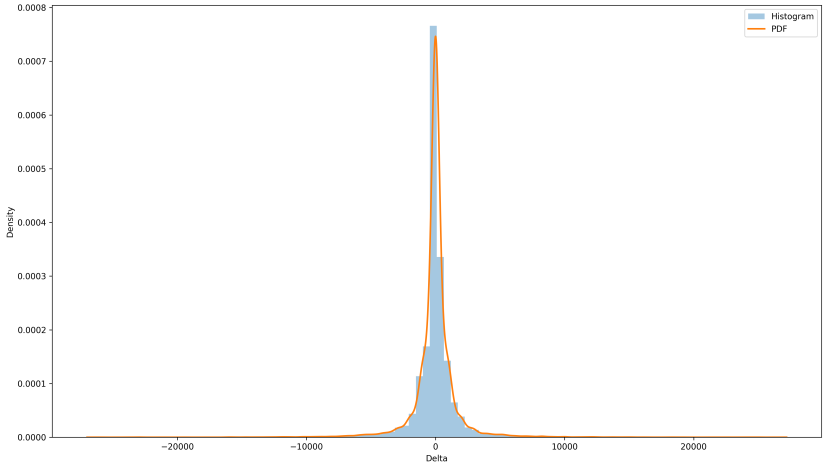

Figure D.12: Zero-truncated delta distribution with PDF of **block 12**

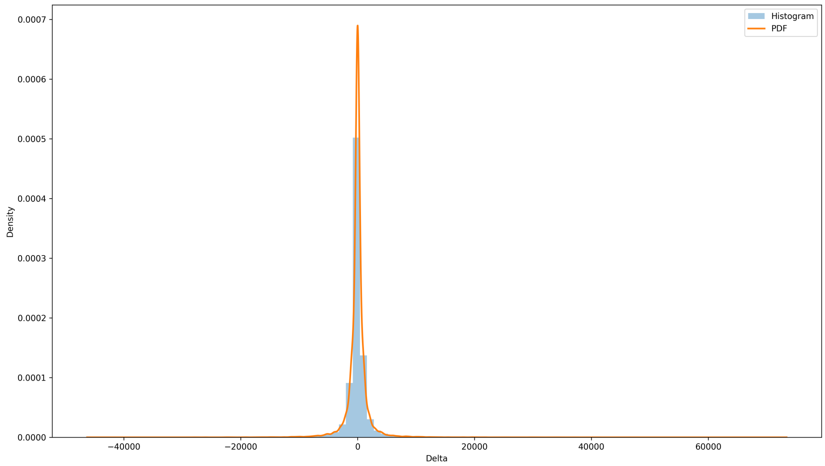

Figure D.13: Zero-truncated delta distribution with PDF of **block 13**

E
Delta distributions of CMG subcategories

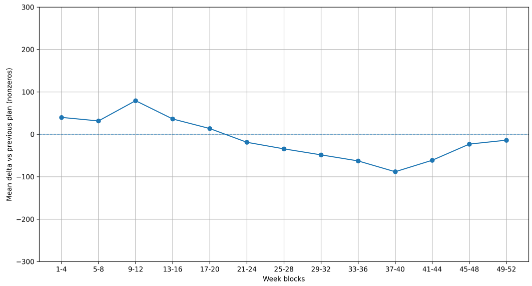

(a) GA02

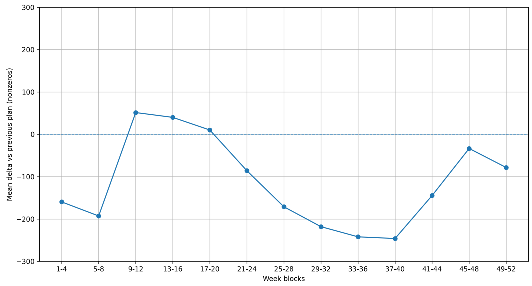

(b) FA01

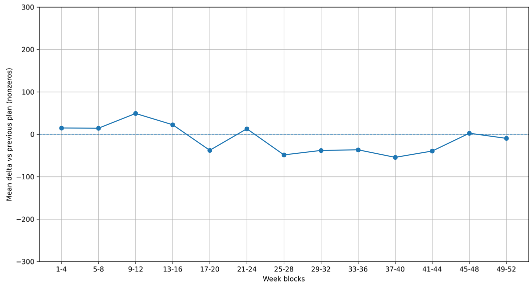

(c) GA03

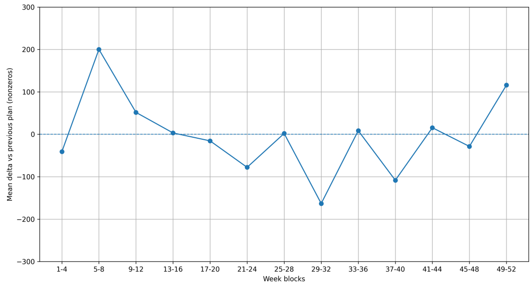

(d) FA02

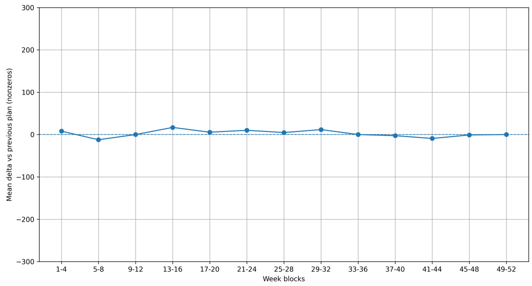

(e) GA11

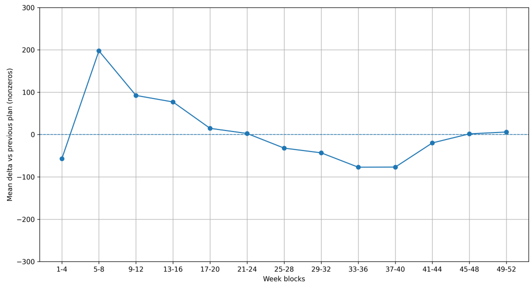

(f) GA07

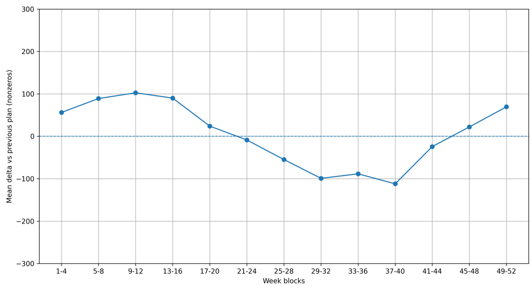

(g) GA06

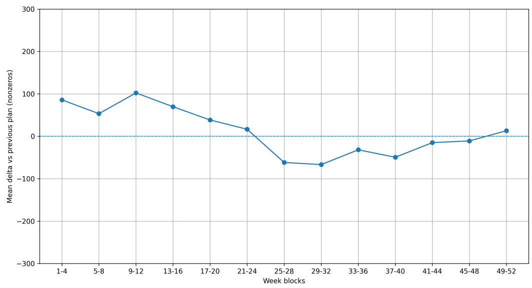

(h) GA10

Figure E.1: Mean of **nonzero** deltas per rolling 4-week block for **subcategories**

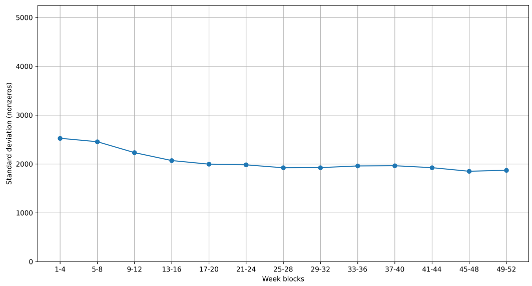

(a) GA02

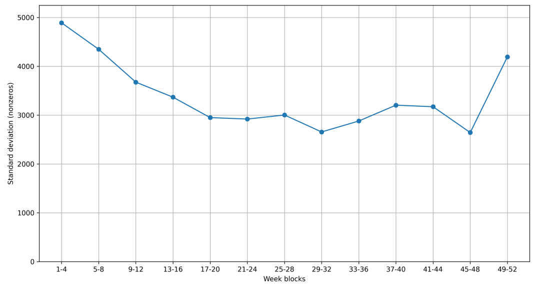

(b) FA01

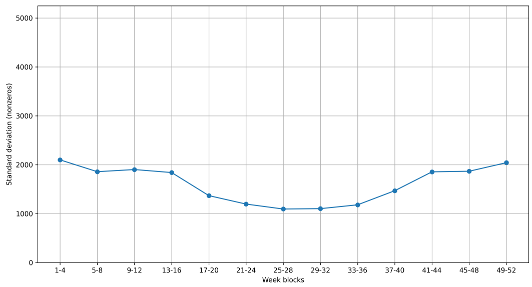

(c) GA03

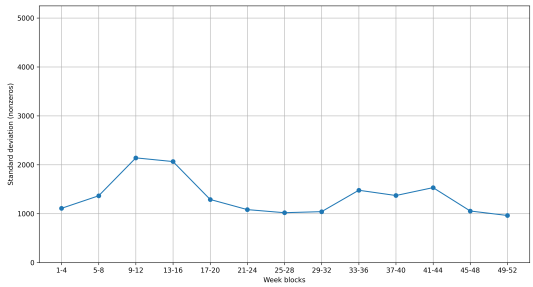

(d) FA02

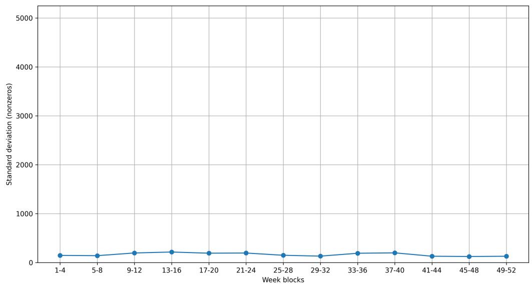

(e) GA11

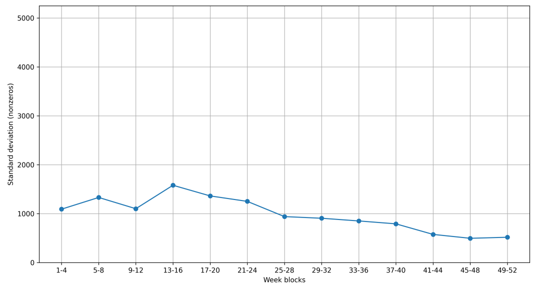

(f) GA07

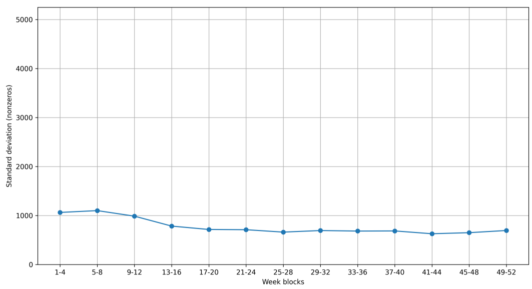

(g) GA06

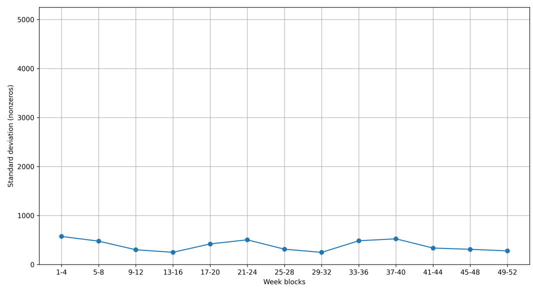

(h) GA10

Figure E.2: Standard deviation of **nonzero** deltas per rolling 4-week block for **subcategories**

F
Demand pattern classification

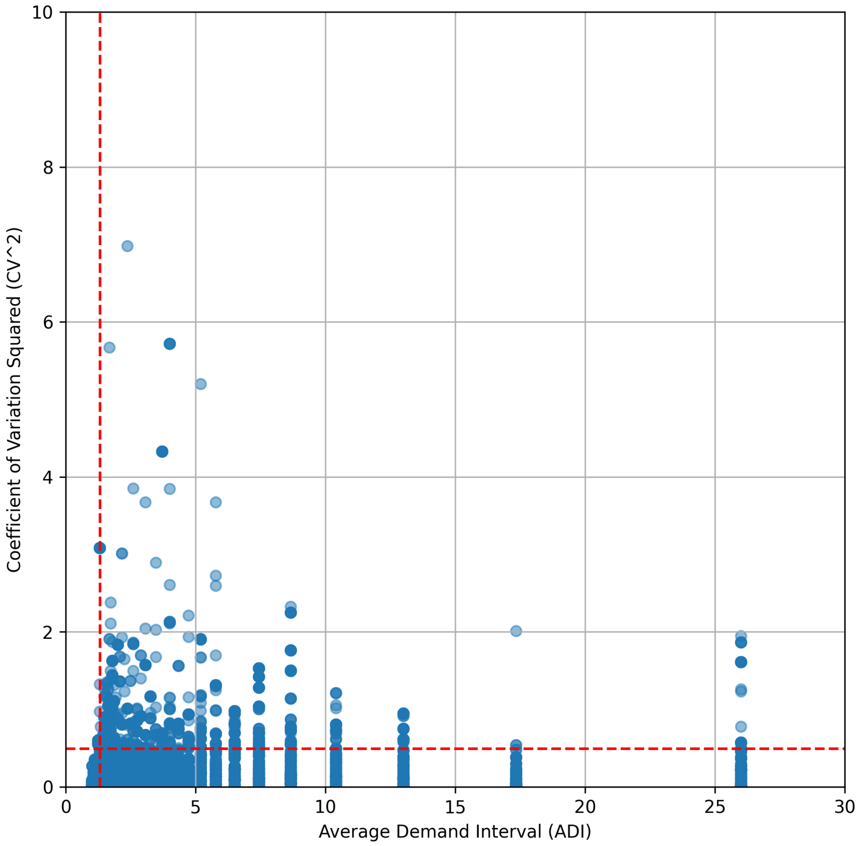

Figure F.1: Material demand classification

G

S2 Number of simulations result

Table G.1: S2 Experiment result

Simulation [–]	Obsolete costs [€]	Sample mean [€]	s [€]	Relative error [-]
1	158 039	158 039	-	-
2	192 955	175 497	24 689	0.195
3	147 834	166 276	23 661	0.161
4	141 051	159 970	23 072	0.141
5	**179 948**	**163 965**	**21 888**	**0.117**
6	148 381	161 368	20 585	0.102
7	176 475	163 526	19 640	0.089
8	161 489	163 272	18 197	0.077
9	147 428	161 511	17 822	0.072
10	**155 190**	**160 879**	**16 921**	**0.065**
11	164 916	161 246	16 099	0.059
12	134 284	158 999	17 210	0.061
13	128 684	156 667	18 499	0.064
14	178 244	158 208	18 685	0.062
15	**160 505**	**158 362**	**18 015**	**0.058**
16	144 194	157 476	17 761	0.055
17	178 003	158 684	17 903	0.054
18	261 491	164 395	29 814	0.084
19	145 400	163 395	29 300	0.081
20	152 358	162 843	28 625	0.077
21	175 490	163 446	28 036	0.073
22	166 865	163 601	27 370	0.070
23	142 972	162 704	27 085	0.068
24	108 801	160 458	28 684	0.072
25	155 416	160 257	28 098	0.069
26	155 449	160 072	27 546	0.066
27	128 726	158 911	27 677	0.066
28	164 072	159 095	27 177	0.063
29	149 168	158 753	26 751	0.061
30	**156 543**	**158 679**	**26 289**	**0.059**

H
A Experiments results

Table H.1: Experiment **A1** result

Iteration (.simulation) [–]	**Obsolete material costs** [€]	**Obsolete material stock** [pieces]	**HSS** [DOS]
1.1	3 180	15 735	0

Table H.2: Experiment **A2** result

Iteration (.simulation) [–]	**Obsolete material costs** [€]	**Obsolete material stock** [pieces]	**HSS** [DOS]
1.1	**28 612**	**330 462**	**1**
2.1	51 180	632 410	2
3.1	73 814	931 612	3
4.1	96 381	1 233 560	4
5.1	119 013	1 536 169	5
6.1	142 184	1 840 156	6
7.1	166 984	2 152 547	7

Table H.3: Experiment **A3** result

Iteration (.simulation) [–]	**Obsolete material costs** [€]	**Obsolete material stock** [pieces]	**HSS** [DOS]
1.1	**24 321**	**308 553**	**1**
2.1	44 794	598 672	2
3.1	65 334	886 045	3
4.1	85 807	1 176 164	4
5.1	106 375	1 467 304	5
6.1	127 422	1 759 102	6
7.1	147 930	2 049 584	7

Table H.4: Experiment **A4** result

Iteration (.simulation) [–]	**Obsolete material costs** [€]	**Obsolete material stock** [pieces]	**HSS** [DOS]
1.1	**24 905**	**312 046**	**1**
2.1	45 902	604 938	2
3.1	67 026	895 804	3
4.1	88 023	1 188 696	4
5.1	109 050	1 481 948	5
6.1	130 745	1 777 900	6
7.1	151 772	2 071 152	7

I

B Experiments results

Table I.1: Experiment **B1** result

Iteration (.simulation) [–]	**Obsolete material costs** [€]	**Obsolete material stock** [pieces]	**HSS** [DOS]
1.1	33 443	183 224	1
1.2	128 619	1 416 651	1
1.3	170 255	2 369 178	1
1.4	26 918	1 849 907	1
1.5	109 729	972 376	1
1.6	1 911	12 163	1
1.7	93 889	868 433	1
1.8	967	6 726	1
1.9	2 770	17 266	1
1.10	131 935	2 318 623	1
1.M	70 044	–	1

Table I.2: Experiment **B2** result **part 1**

Iteration (.simulation) [–]	Obsolete material costs [€]	Obsolete material stock [pieces]	HSS [DOS]
1.1	57 277	485 637	1
1.2	155 299	1 689 561	1
1.3	195 615	2 646 953	1
1.4	52 198	2 106 198	1
1.5	125 218	1 209 495	1
1.6	24 347	313 545	1
1.7	116 319	1 164 342	1
1.8	22 724	301 600	1
1.9	25 069	315 604	1
1.10	149 215	2 503 879	1
1.M	**92 328**	**–**	**1**
2.1	78 062	775 716	2
2.2	179 567	1 952 366	2
2.3	208 445	2 807 588	2
2.4	74 614	2 349 712	2
2.5	152 683	1 548 957	2
2.6	45 997	606 435	2
2.7	138 140	1 458 096	2
2.8	45 771	604 792	2
2.9	46 096	608 313	2
2.10	166 220	2 689 975	2
2.M	113 560	–	2
3.1	100 300	1 073 495	3
3.2	193 964	2 125 497	3
3.3	234 465	3 076 363	3
3.4	97 697	2 595 926	3
3.5	180 418	1 885 219	3
3.6	70 620	917 772	3
3.7	160 530	1 753 542	3
3.8	68 178	905 413	3
3.9	70 806	921 291	3
3.10	183 769	2 876 671	3
3.M	136 075	–	3

Table I.3: Experiment **B2** result **part 2**

Iteration (.simulation) [–]	**Obsolete material costs** [€]	**Obsolete material stock** [pieces]	**HSS** [DOS]
4.1	121 084	1 363 574	4
4.2	218 200	2 387 942	4
4.3	248 494	3 249 507	4
4.4	120 113	2 839 440	4
4.5	195 115	2 119 981	4
4.6	93 143	1 219 293	4
4.7	182 352	2 047 358	4
4.8	91 197	1 208 476	4
4.9	92 102	1 215 166	4
4.10	201 453	3 065 826	4
4.M	156 325	–	4
5.1	143 192	1 663 193	5
5.2	234 430	2 578 627	5
5.3	272 545	3 511 245	5
5.4	143 823	3 092 134	5
5.5	223 468	2 463 163	5
5.6	115 394	1 518 506	5
5.7	204 074	2 340 042	5
5.8	113 576	1 508 968	5
5.9	114 280	1 511 655	5
5.10	219 017	3 252 681	5
5.M	178 380	–	5
6.1	164 131	1 952 466	6
6.2	256 993	2 826 126	6
6.3	299 019	3 790 162	6
6.4	165 638	3 330 202	6
6.5	238 507	2 700 404	6
6.6	137 721	1 817 723	6
6.7	225 955	2 630 965	6
6.8	136 034	1 806 771	6
6.9	137 636	1 817 859	6
6.10	236 100	3 436 390	6
6.M	199 773	–	6
7.1	187 739	2 259 105	7
7.2	285 437	3 148 376	7
7.3	311 906	3 954 677	7
7.4	191 544	3 592 976	7
7.5	263 823	3 022 661	7
7.6	158 796	2 108 989	7
7.7	247 714	2 924 156	7
7.8	158 401	2 107 206	7
7.9	160 633	2 120 166	7
7.10	254 004	3 625 027	7
7.M	222 000	–	7

Table I.4: Experiment **B3** result **part 1**

Iteration (.simulation) [–]	**Obsolete material costs** [€]	**Obsolete material stock** [pieces]	**HSS** [DOS]
1.1	52 986	463 728	1
1.2	149 805	1 656 760	1
1.3	192 318	2 624 232	1
1.4	47 907	2 084 289	1
1.5	121 921	1 186 774	1
1.6	22 253	301 716	1
1.7	113 022	1 141 621	1
1.8	20 630	289 771	1
1.9	22 835	303 135	1
1.10	145 779	2 480 517	1
1.M	**88 946**	**–**	**1**
2.1	71 676	741 978	2
2.2	170 775	1 896 844	2
2.3	201 851	2 762 146	2
2.4	68 228	2 315 974	2
2.5	146 120	1 503 875	2
2.6	41 809	582 777	2
2.7	131 545	1 412 654	2
2.8	41 582	581 134	2
2.9	41 877	584 295	2
2.10	159 626	2 644 533	2
2.M	107 509	–	2
3.1	91 820	1 027 928	3
3.2	181 875	2 047 254	3
3.3	224 574	3 008 200	3
3.4	89 217	2 550 359	3
3.5	170 526	1 817 056	3
3.6	64 337	882 285	3
3.7	150 639	1 685 379	3
3.8	61 896	869 926	3
3.9	64 524	885 804	3
3.10	173 878	2 808 508	3
3.M	127 329	–	3

Table I.5: Experiment **B3** result **part 2**

Iteration (.simulation) [–]	**Obsolete material costs** [€]	**Obsolete material stock** [pieces]	**HSS** [DOS]
4.1	110 510	1 306 178	4
4.2	202 815	2 286 978	4
4.3	235 306	3 158 623	4
4.4	109 539	2 782 044	4
4.5	181 927	2 029 097	4
4.6	84 766	1 171 977	4
4.7	169 164	1 956 474	4
4.8	82 820	1 161 160	4
4.9	83 725	1 167 850	4
4.10	188 265	2 974 942	4
4.M	144 884	–	4
5.1	130 493	1 593 608	5
5.2	215 747	2 454 942	5
5.3	256 060	3 397 640	5
5.4	131 155	3 022 909	5
5.5	206 983	2 349 558	5
5.6	104 923	1 459 361	5
5.7	187 588	2 226 437	5
5.8	103 106	1 449 823	5
5.9	103 810	1 452 510	5
5.10	202 532	3 139 076	5
5.M	164 240	–	5
6.1	149 369	1 871 412	6
6.2	235 614	2 685 166	6
6.3	279 051	3 654 282	6
6.4	150 875	3 249 148	6
6.5	219 357	2 569 884	6
6.6	125 156	1 746 749	6
6.7	206 774	2 500 085	6
6.8	123 469	1 735 797	6
6.9	125 041	1 746 525	6
6.10	216 919	3 305 510	6
6.M	183 163	–	6
7.1	168 685	2 156 142	7
7.2	258 564	2 974 615	7
7.3	288 640	3 796 076	7
7.4	172 490	3 490 013	7
7.5	241 345	2 869 060	7
7.6	144 137	2 026 186	7
7.7	225 236	2 770 555	7
7.8	143 742	2 024 403	7
7.9	145 835	2 036 723	7
7.10	231 387	3 470 785	7
7.M	202 006	–	7

Table I.6: Experiment **B4** result **part 1**

Iteration (.simulation) [–]	Obsolete material costs [€]	Obsolete material stock [pieces]	HSS [DOS]
1.1	52 986	463 728	1
1.2	149 488	1 653 037	1
1.3	182 068	2 536 236	1
1.4	47 907	2 084 289	1
1.5	121 487	1 184 208	1
1.6	22 155	300 575	1
1.7	113 022	1 141 621	1
1.8	20 630	289 771	1
1.9	21 566	295 446	1
1.10	63 356	1 512 837	1
1.M	**79 467**	**–**	**1**
2.1	71 676	741 978	2
2.2	162 936	1 827 069	2
2.3	198 378	2 731 317	2
2.4	68 228	2 315 974	2
2.5	140 417	1 457 721	2
2.6	41 787	582 517	2
2.7	131 527	1 412 510	2
2.8	41 582	581 134	2
2.9	41 399	581 645	2
2.10	81 112	1 722 774	2
2.M	97 904	–	2
3.1	91 851	1 028 288	3
3.2	180 741	2 037 304	3
3.3	215 786	2 934 524	3
3.4	89 248	2 550 719	3
3.5	161 569	1 740 415	3
3.6	64 337	882 285	3
3.7	150 639	1 685 379	3
3.8	61 896	869 926	3
3.9	62 716	875 113	3
3.10	99 795	1 938 748	3
3.M	117 858	–	3

Table I.7: Experiment **B4** result **part 2**

Iteration (.simulation) [–]	**Obsolete material costs** [€]	**Obsolete material stock** [pieces]	**HSS** [DOS]
4.1	110 510	1 306 178	4
4.2	197 368	2 238 833	4
4.3	233 500	3 144 180	4
4.4	109 539	2 782 044	4
4.5	179 951	2 011 689	4
4.6	84 766	1 171 977	4
4.7	169 113	1 955 908	4
4.8	82 820	1 161 160	4
4.9	82 678	1 161 838	4
4.10	118 107	2 151 262	4
4.M	136 835	–	4
5.1	130 039	1 588 953	5
5.2	215 294	2 450 287	5
5.3	251 029	3 353 070	5
5.4	130 670	3 017 894	5
5.5	199 722	2 289 663	5
5.6	104 923	1 459 361	5
5.7	187 588	2 226 437	5
5.8	103 106	1 449 823	5
5.9	103 432	1 450 097	5
5.10	136 789	2 367 236	5
5.M	156 259	–	5
6.1	149 212	1 869 786	6
6.2	232 247	2 655 205	6
6.3	268 079	3 560 385	6
6.4	150 718	3 247 522	6
6.5	218 040	2 559 937	6
6.6	125 156	1 746 749	6
6.7	206 743	2 499 725	6
6.8	123 469	1 735 797	6
6.9	123 394	1 736 822	6
6.10	155 101	2 579 750	6
6.M	175 216	–	6
7.1	168 260	2 151 785	7
7.2	250 585	2 903 368	7
7.3	285 277	3 767 892	7
7.4	172 065	3 485 656	7
7.5	237 811	2 837 911	7
7.6	144 137	2 026 186	7
7.7	225 236	2 770 555	7
7.8	143 742	2 024 403	7
7.9	144 919	2 031 308	7
7.10	231 356	3 470 425	7
7.M	200 339	–	7

J
C Experiments results

Table J.1: Experiment **C1** result

Iteration (.simulation) [–]	**Obsolete material costs** [€]	**Obsolete material stock** [pieces]	**HSS** [DOS]
1.1	2 920	13 460	0

Table J.2: Experiment **C2** result

Iteration (.simulation) [–]	**Obsolete material costs** [€]	**Obsolete material stock** [pieces]	**HSS** [DOS]
1.1	**28 240**	**327 171**	**1**
2.1	50 695	628 103	2
3.1	73 818	931 735	3
4.1	96 274	1 232 667	4
5.1	118 793	1 534 260	5
6.1	141 852	1 837 231	6
7.1	166 540	2 148 606	7

Table J.3: Experiment **C3** result

Iteration (.simulation) [–]	**Obsolete material costs** [€]	**Obsolete material stock** [pieces]	**HSS** [DOS]
1.1	**23 948**	**305 262**	**1**
2.1	44 310	594 365	2
3.1	65 338	886 168	3
4.1	85 700	1 175 271	4
5.1	106 125	1 465 035	5
6.1	127 089	1 756 177	6
7.1	147 486	2 045 643	7

Table J.4: Experiment **C4** result

Iteration (.simulation) [–]	**Obsolete material costs** [€]	**Obsolete material stock** [pieces]	**HSS** [DOS]
1.1	**24 471**	**308 035**	**1**
2.1	45 417	600 631	2
3.1	67 000	895 567	3
4.1	87 915	1 187 803	4
5.1	108 830	1 480 039	5
6.1	130 412	1 774 975	6
7.1	151 328	2 067 211	7

K
D Experiments results

Table K.1: Experiment **D1** result

Iteration (.simulation) [–]	**Obsolete material costs** [€]	**Obsolete material stock** [pieces]	**HSS** [DOS]
1.1	2 889	13 100	0

Table K.2: Experiment **D2** result

Iteration (.simulation) [–]	**Obsolete material costs** [€]	**Obsolete material stock** [pieces]	**HSS** [DOS]
1.1	**28 209**	**326 811**	**1**
2.1	50 665	627 743	2
3.1	73 787	931 375	3
4.1	96 245	1 232 328	4
5.1	118 762	1 533 900	5
6.1	141 821	1 836 871	6
7.1	166 509	2 148 246	7

Table K.3: Experiment **D3** result

Iteration (.simulation) [–]	**Obsolete material costs** [€]	**Obsolete material stock** [pieces]	**HSS** [DOS]
1.1	**23 918**	**304 902**	**1**
2.1	44 279	594 005	2
3.1	65 307	885 808	3
4.1	85 669	1 174 911	4
5.1	106 094	1 464 675	5
6.1	127 059	1 755 817	6
7.1	147 455	2 045 283	7

Table K.4: Experiment **D4** result

Iteration (.simulation) [–]	**Obsolete material costs** [€]	**Obsolete material stock** [pieces]	**HSS** [DOS]
1.1	**24 502**	**308 395**	**1**
2.1	45 387	600 271	2
3.1	67 000	895 567	3
4.1	87 884	1 187 443	4
5.1	108 799	1 479 679	5
6.1	130 382	1 774 615	6
7.1	151 297	2 066 851	7

L

Perturbation size Experiments results

Table L.1: **Perturbation size** Experiment **A1** item-level result

Iteration.simulation [–]	**Item** [–]	**Obsolete material costs** [€]	**Obsolete material stock** [pieces]	**HSS** [DOS]
	1	0	0	0
	2	291	2 635	0
	3	0	0	0
	4	0	0	0
	5	284	1 151	0
1.1	total	**575**	**3 786**	**0**

Table L.2: **Perturbation size** Experiment **A4** item-level result

Iteration.simulation [–]	Item [–]	Obsolete material costs [€]	Obsolete material stock [pieces]	HSS [DOS]
	1	4 179	49 063	
	2	1 494	13 527	
	3	378	58 178	
	4	657	10 084	
	5	1 619	6 551	
1.1	total	**8 327**	**137 403**	**1**
	1	8 297	97 406	
	2	2 697	24 419	
	3	756	116 356	
	4	1 315	20 168	
	5	2 286	9 251	
2.1	total	15 350	267 600	2
	1	12 476	146 469	
	2	3 298	29 865	
	3	1 134	174 534	
	4	1 972	30 252	
	5	3 620	14 651	
3.1	total	22 500	395 771	3
	1	16 593	194 812	
	2	4 501	40 757	
	3	1 512	232 712	
	4	2 630	40 336	
	5	4 287	17 351	
4.1	total	29 523	525 968	4
	1	20 742	243 515	
	2	5 704	51 649	
	3	1 890	290 890	
	4	3 287	50 420	
	5	4 954	20 051	
5.1	total	36 577	656 525	5
	1	24 890	292 218	
	2	6 907	62 541	
	3	2 268	349 068	
	4	3 944	60 504	
	5	6 288	25 451	
6.1	total	44 298	789 782	6
	1	29 038	340 921	
	2	8 110	73 433	
	3	2 646	407 246	
	4	4 602	70 588	
	5	6 955	28 151	
7.1	total	51 351	920 339	7

Table L.3: **Perturbation size** Experiment **D1** item-level result **(P = 123)**

Iteration.simulation [–]	**Item** [–]	**Obsolete material costs** [€]	**Obsolete material stock** [pieces]	**HSS** [DOS]
	1	31	360	0
	2	0	0	0
	3	0	0	0
	4	0	0	0
	5	566	2 290	0
1.1	total	**596**	**2 650**	0

Table L.4: **Perturbation size** Experiment **D1** item-level result **(P = 1234)**

Iteration.simulation [–]	**Item** [–]	**Obsolete material costs** [€]	**Obsolete material stock** [pieces]	**HSS** [DOS]
	1	0	0	0
	2	0	0	0
	3	0	0	0
	4	0	0	0
	5	418	1 692	0
1.1	total	**418**	**1 692**	0

Table L.5: **Perturbation size** Experiment **D1** item-level result **(P = 12345)**

Iteration.simulation [–]	**Item** [–]	**Obsolete material costs** [€]	**Obsolete material stock** [pieces]	**HSS** [DOS]
	1	0	0	0
	2	0	0	0
	3	0	0	0
	4	0	0	0
	5	0	0	0
1.1	total	**0**	**0**	0

Table L.6: **Perturbation size** Experiment **D1** item-level result **(P = 123456)**

Iteration.simulation [–]	**Item** [–]	**Obsolete material costs** [€]	**Obsolete material stock** [pieces]	**HSS** [DOS]
	1	0	0	0
	2	0	0	0
	3	0	0	0
	4	0	0	0
	5	0	0	0
1.1	total	**0**	**0**	0

Table L.7: Perturbation size Experiment **D4** item-level result **(P = 123)**

Iteration.simulation [–]	**Item** [–]	**Obsolete material costs** [€]	**Obsolete material stock** [pieces]	**HSS** [DOS]
	1	4 148	48 703	
	2	1 091	9 876	
	3	378	58 178	
	4	657	10 084	
	5	1 233	4 990	
1.1	total	**7 507**	**131 831**	**1**
	1	8 297	97 406	
	2	2 181	19 752	
	3	756	116 356	
	4	1 315	20 168	
	5	2 567	10 390	
2.1	total	15 116	264 072	2
	1	12 476	146 469	
	2	3 272	29 628	
	3	1 134	174 534	
	4	1 972	30 252	
	5	3 234	13 090	
3.1	total	22 088	393 973	3
	1	16 593	194 812	
	2	4 363	39 504	
	3	1 512	232 712	
	4	2 630	40 336	
	5	3 957	16 016	
4.1	total	29 055	523 380	4
	1	20 772	243 875	
	2	5 454	49 380	
	3	1 890	290 890	
	4	3 287	50 420	
	5	5 235	21 190	
5.1	total	36 638	655 755	5
	1	24 890	292 218	
	2	6 544	59 256	
	3	2 268	349 068	
	4	3 944	60 504	
	5	5 936	24 024	
6.1	total	43 582	785 070	6
	1	29 069	341 281	
	2	7 635	69 132	
	3	2 646	407 246	
	4	4 602	70 588	
	5	7 237	29 290	
7.1	total	51 188	917 537	7

Table L.8: Perturbation size Experiment **D4** item-level result **(P = 1234)**

Iteration.simulation [–]	Item [–]	Obsolete material costs [€]	Obsolete material stock [pieces]	HSS [DOS]
	1	4 148	48 703	
	2	1 091	9 876	
	3	378	58 178	
	4	657	10 084	
	5	1 085	4 392	
1.1	total	**7 360**	**131 233**	1
	1	8 297	97 406	
	2	2 181	19 752	
	3	756	116 356	
	4	1 315	20 168	
	5	2 419	9 792	
2.1	total	14 968	263 474	2
	1	12 445	146 109	
	2	3 272	29 628	
	3	1 134	174 534	
	4	1 972	30 252	
	5	3 086	12 492	
3.1	total	21 910	393 015	3
	1	16 593	194 812	
	2	4 363	39 504	
	3	1 512	232 712	
	4	2 630	40 336	
	5	4 421	17 892	
4.1	total	29 518	525 256	4
	1	20 742	243 515	
	2	5 454	49 380	
	3	1 890	290 890	
	4	3 287	50 420	
	5	5 088	20 592	
5.1	total	36 460	654 797	5
	1	24 890	292 218	
	2	6 544	59 256	
	3	2 268	349 068	
	4	3 944	60 504	
	5	6 422	25 992	
6.1	total	44 068	787 038	6
	1	29 038	340 921	
	2	7 635	69 132	
	3	2 646	407 246	
	4	4 602	70 588	
	5	7 089	28 692	
7.1	total	51 010	916 579	7

Table L.9: Perturbation size Experiment **D4** item-level result **(P = 12345)**

Iteration.simulation [–]	**Item** [–]	**Obsolete material costs** [€]	**Obsolete material stock** [pieces]	**HSS** [DOS]
	1	4 148	48 703	
	2	1 091	9 876	
	3	378	58 178	
	4	657	10 084	
	5	989	4 004	
1.1	total	7 264	130 845	1
	1	8 297	97 406	
	2	2 181	19 752	
	3	756	116 356	
	4	1 315	20 168	
	5	1 979	8 008	
2.1	total	14 527	261 690	2
	1	12 445	146 109	
	2	3 272	29 628	
	3	1 134	174 534	
	4	1 972	30 252	
	5	2 968	12 012	
3.1	total	21 791	392 535	3
	1	16 593	194 812	
	2	4 363	39 504	
	3	1 512	232 712	
	4	2 630	40 336	
	5	3 957	16 016	
4.1	total	29 055	523 380	4
	1	20 742	243 515	
	2	5 454	49 380	
	3	1 890	290 890	
	4	3 287	50 420	
	5	4 946	20 020	
5.1	total	36 318	654 225	5
	1	24 890	292 218	
	2	6 544	59 256	
	3	2 268	349 068	
	4	3 944	60 504	
	5	5 936	24 024	
6.1	total	43 582	785 070	6
	1	29 038	340 921	
	2	7 635	69 132	
	3	2 646	407 246	
	4	4 602	70 588	
	5	6 925	28 028	
7.1	total	50 846	915 915	7

Table L.10: Perturbation size Experiment **D4** item-level result **(P = 123456)**

Iteration.simulation [–]	Item [–]	Obsolete material costs [€]	Obsolete material stock [pieces]	HSS [DOS]
	1	4 148	48 703	
	2	1 091	9 876	
	3	378	58 178	
	4	657	10 084	
	5	989	4 004	
1.1	total	7 264	130 845	1
	1	8 297	97 406	
	2	2 181	19 752	
	3	756	116 356	
	4	1 315	20 168	
	5	1 979	8 008	
2.1	total	14 527	261 690	2
	1	12 445	146 109	
	2	3 272	29 628	
	3	1 134	174 534	
	4	1 972	30 252	
	5	2 968	12 012	
3.1	total	21 791	392 535	3
	1	16 593	194 812	
	2	4 363	39 504	
	3	1 512	232 712	
	4	2 630	40 336	
	5	3 957	16 016	
4.1	total	29 055	523 380	4
	1	20 742	243 515	
	2	5 454	49 380	
	3	1 890	290 890	
	4	3 287	50 420	
	5	4 946	20 020	
5.1	total	36 318	654 225	5
	1	24 890	292 218	
	2	6 544	59 256	
	3	2 268	349 068	
	4	3 944	60 504	
	5	5 936	24 024	
6.1	total	43 582	785 070	6
	1	29 038	340 921	
	2	7 635	69 132	
	3	2 646	407 246	
	4	4 602	70 588	
	5	6 925	28 028	
7.1	total	50 846	915 915	7

M
Scientific paper

Optimising safety stock levels of changeover materials using simulation-based optimisation

Gijs Janssen

Faculty of Civil Engineering and Geosciences
Delft University of Technology, Delft, Netherlands

ABSTRACT

Material Requirements Planning (MRP) is widely used to translate production plans into procurement and inventory decisions, yet its performance deteriorates in environments characterised by demand uncertainty and frequent material changeovers. In such contexts, deterministic MRP models often reduce safety stock levels of changeover materials to zero in order to minimise obsolete inventory, thereby increasing vulnerability to stockouts, minimum-order-quantity rush orders, and last-minute procurement costs. This study investigates whether simulation-based optimisation can be used to parametrise safety stock levels for changeover materials in a deterministic mixed-integer programming MRP model so as to improve rolling-horizon performance in a stochastic production environment. A simulation-based optimisation framework is developed in which the manufacturer's existing MRP optimisation tool is treated as a black box and embedded in an Optimisation with Simulation-based Iterations structure. To better represent robustness in the planning logic, a hard safety stock constraint is introduced and safety stock levels are optimised for grouped materials. Demand uncertainty is modelled using historical weekly production plans and analysed through block bootstrapping and kernel density estimation. The results show that the proposed framework is technically feasible and provides a reproducible method for evaluating safety stock policies under uncertainty. However, the final optimisation does not improve system-wide obsolete stock costs compared with the baseline, and low safety stock levels consistently perform best. The main contribution of this study is therefore methodological: it demonstrates how simulation-based optimisation can be applied in an industry-scale, black-box MRP environment while also highlighting the limitations of optimising robustness through obsolete stock costs alone.

Keywords: Material Requirements Planning; Safety Stock; Changeover Materials; Simulation-Based Optimisation; Simulated Annealing; Demand Uncertainty; Rolling Horizon Planning.

1. INTRODUCTION

Material Requirements Planning (MRP) is a core mechanism in manufacturing supply chains for translating production schedules into purchasing and inventory decisions [1]. Its role is particularly important in industries where material costs constitute a large share of total production costs and where production continuity depends on timely replenishment [2]. However, deterministic MRP logic becomes less effective when production environments are exposed to forecast error, changing customer demand, and product or packaging changeovers. In such settings, planning decisions that appear cost-efficient in a deterministic run may perform poorly over time when demand deviates from expectations. This problem is especially relevant for changeover materials. These materials are required when a production line switches from one product variant or packaging format to another. If the changeover takes place too early, leftover materials may become obsolete; if it takes place too late, production may face stockouts, emergency orders, or missed campaign windows. In industrial practice, safety stocks for changeover materials are often minimised or set to zero to reduce the risk of write-offs. Although that policy may be efficient in a deterministic setting, it can create high vulnerability in a stochastic planning environment.

At the same time, advanced industrial MRP tools rely on deterministic mixed-integer programming models that are solved repeatedly in a rolling-horizon manner. These models support tactical decision-making, yet robustness against uncertainty is still often handled through static planning parameters such as safety stocks. This creates a gap between deterministic optimisation and the stochastic reality of sequential planning. Viewed from a sequential decision perspective, the performance of MRP should not be assessed through a single run, but rather through repeated decisions over time under uncertainty.

This study addresses that gap by asking: how can simulation-based optimisation be used to parametrise safety stocks of changeover materials in a deterministic MIP model such that rolling-horizon performance improves in a stochastic MRP environment? To answer this question, a simulation-based optimisation framework is developed around an existing industrial planning tool, with the objective of learning safety stock policies that are robust over repeated planning cycles.

2. LITERATURE REVIEW

Material Requirements Planning (MRP) remains one of the most widely used approaches for translating production schedules into procurement and inventory decisions in manufacturing environments. Although MRP has demonstrated strong stability when compared with alternative production control systems such as Kanban, CONWIP, and Drum-Buffer-Rope (DBR), its performance is highly dependent on the proper selection of planning parameters and its ability to cope with environmental uncertainty [3]. In practice, deterministic MRP systems are frequently re-optimised in rolling-horizon settings, where updated information continuously changes inventory and procurement decisions. As a result, uncertainty management has become a central topic in both production planning and supply chain research.

2.1. MRP under uncertainty

Uncertainty in supply chains can originate from multiple sources, including demand variability, lead time fluctuations, capacity limitations, and yield uncertainty. However, the literature consistently identifies demand uncertainty as the most frequently studied source of disruption in supply chain planning and network design [4], [5]. Demand uncertainty arises from forecasting errors, variability in customer orders, and changes in demand timing, all of which can lead to material shortages or excess inventory [5], [6].

To mitigate uncertainty, organisations typically introduce buffers into their planning systems. Safety stocks and safety lead times are among the most commonly applied approaches for protecting against demand and lead time variability, while supplier diversification is frequently proposed as a mitigation strategy for yield uncertainty [5], [7]. Nevertheless, recent research increasingly argues that static planning parameters are insufficient in dynamic environments where forecasts are periodically updated and uncertainty evolves over time [8]. Consequently, there is growing interest in developing adaptive approaches for determining MRP planning parameters.

2.2. Parametrisation of planning parameters

Traditional MRP implementations often determine planning parameters independently, resulting in sub-optimal planning decisions when uncertainty is present. Thevenin et al. demonstrated that jointly optimising safety stocks and lot sizes through stochastic optimisation yields substantially better results than treating both decisions separately [9].

Gansterer et al. found that a simulation-based optimisation approach can effectively improve inventory and service-level performance while simultaneously evaluating alternative planning parameter configurations [10]. As industrial problems become increasingly high-dimensional, additional techniques are required to maintain computational tractability. Werth et al. addressed this challenge by combining simulation-based optimisation with metaheuristic search and material clustering techniques. This study demonstrates that grouping materials according to technological or production characteristics can significantly improve optimisation efficiency [11].

Additional work by Seiringer et al. further demonstrated that simulation-based parameter optimisation can identify planning parameter configurations that better absorb uncertainty than traditional manually configured settings [12], [13].

2.3. Simulation-based optimisation

Simulation-based optimisation (Sim-Opt) has emerged as a promising methodology for improving decision-making in systems characterised by uncertainty and analytical complexity. Sim-Opt integrates simulation and optimisation techniques to evaluate candidate solutions under stochastic conditions [14]. According to Figueira and Almada-Lobo, simulation can be used for solution evaluation, surrogate model construction, analytical model enhancement, or solution generation, depending on the structure of the optimisation problem [14].

A particularly relevant Sim-Opt structure for MRP applications is Optimisation with Simulation-based Iterations (OSI), where simulation is embedded within the optimisation process and repeatedly evaluates candidate solutions under uncertainty. This structure is especially suitable for environments characterised by stochastic behaviour, non-linear interactions and complex decision dependencies, although it can become computationally expensive because many simulation runs are required throughout the optimisation process [14].

A wide range of optimisation techniques has been applied within simulation-based optimisation. Recent MRP studies frequently combine simulation with metaheuristic algorithms such as Simulated Annealing, Genetic Algorithms, Particle Swarm Optimisation, Tabu Search, and Variable Neighbourhood Search [12], [15]. Seiringer et al. further demonstrated that combining Simulated Annealing with Simulation Budget Management improves search efficiency by allocating computational effort primarily to promising solutions [12].

2.4. Research gap

The reviewed literature reveals several important gaps. First, current research predominantly focuses on demand quantity uncertainty, while demand timing uncertainty receives considerably less attention [8], [13]. Second, most studies optimise planning parameters for standard materials rather than changeover materials, despite the unique obsolescence risks associated with inventory held during product transitions [5], [16]. Third, few studies investigate simulation-based optimisation in industry-scale black-box MRP environments where the underlying optimisation model cannot be reformulated analytically [10], [11].

3. RESEARCH APPROACH

This study addresses the research gaps by applying simulation-based optimisation to the parametrisation of safety stocks for changeover materials in a large-scale industrial MRP environment. The proposed framework treats the existing deterministic MRP optimisation model as a black box and evaluates safety stock policies within a stochastic rolling-horizon setting. The methodology integrates uncertainty in the simulations to imitate the real sequential decision analytics planning problem.

3.1. Use case

The study is based on the MRP process of a large manufacturing company operating a complex global supply chain. The company uses a tactical planning tool, referred to as **MRP-Opt**, to support material planning and changeover scheduling. The model underlying this tool is a deterministic mixed-integer programming formulation that is solved weekly and incorporates updated case-file data. Because the model is highly complex and cannot be analytically represented in a tractable way, it is treated as a black box in this study.

MRP-Opt optimises changeovers while balancing material stocks, changeover timing, and planning constraints across the supply chain. The model includes multiple prioritised objectives, of which the most important is the minimisation of obsolete stock costs. Other objectives include limiting material slack during changeovers, completing changeovers in time, reducing safety stock slack, and penalising unfavourable ordering patterns. Because obsolete stock cost minimisation has the highest priority, the model often allows stock levels to fall below nominal safety stock targets, especially for changeover materials.

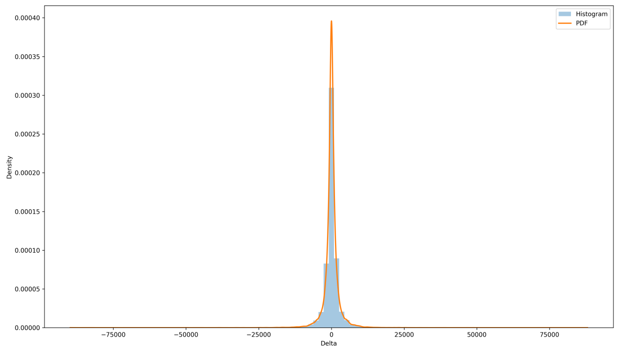

Figure 1. Zero-truncated delta distribution (block 1)

3.2. Data preprocessing

Demand uncertainty is modelled using historical weekly production plans collected over approximately six months. To understand the structure of variability, the production plans are analysed using methods based on **block bootstrapping** and **kernel density estimation** (KDE). The analysis shows that almost all observed variability is demand-driven (99.98%), while planning-induced variability is negligible. Using block bootstrapping, the the production plans are separated into 4-week blocks. The difference in production demand between consecutive production plans (deltas) was analysed for all blocks. The deltas are found to be heavily zero-inflated, with roughly 70–72% zero deltas depending on the planning block. For the nonzero values, zero-truncated distributions are fitted using KDE and are later used in the simulation framework. Figure 1 shows the zero-truncated delta distribution of block 1 as an example.

3.3. Material classification

Changeover materials were classified into groups based on their demand behaviour and economic importance to reduce the dimensionality of the safety stock optimisation problem. It is a way of grouping materials for which a shared safety stock policy can be found. Demand behaviour was characterised using the classification framework proposed by Syntetos et al., which categorises items according to the **Average Inter-Demand Interval** (ADI) and the squared **Coefficient of Variation** (CV^2) [17]. Based on these metrics, materials were classified as smooth, erratic, intermittent, or lumpy.

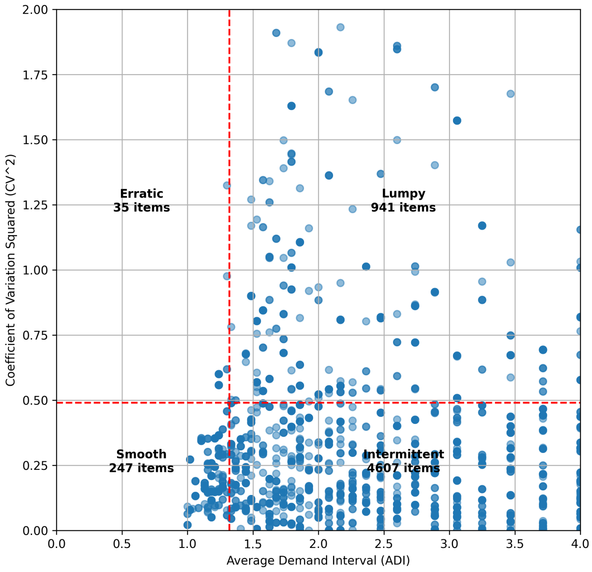

Figure 2. Material demand classification of changeover materials

Approximately 79% of items exhibit intermittent demand. This shows that demand occurs infrequently and is separated by relatively long periods without consumption. However, when demand does occur, the quantity remains relatively stable, as shown by the low CV^2 values in Figure 2. While 16% of materials are classified as lumpy, only a small fraction of materials show smooth or erratic demand behaviour. This suggests that the primary source of uncertainty for most changeover materials is the timing of demand rather than the quantity required.

In addition to demand behaviour, an ABC analysis was performed to distinguish materials according to their contribution to total inventory value. Following the Pareto principle, items were ranked by annual demand value and divided into three classes shown in Figure 3. Class A materials represent the most economically important items, accounting for roughly 80% of total inventory value despite comprising only a small proportion of the portfolio [18]. Class B materials contribute an additional 15% of value, while the remaining Class C materials account for only a minor share of total inventory value [18].

The two classification dimensions were subsequently

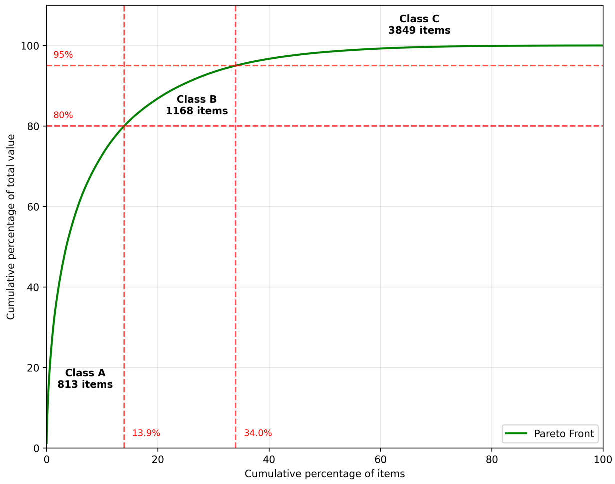

Figure 3. Pareto distribution of ABC analysis

combined to create material groups that capture both operational uncertainty and economic relevance. This resulted in twelve material groups, each representing a unique combination of ABC class and demand pattern. Such grouping enables the optimisation framework to assign a common safety stock parameter to multiple materials while maintaining a meaningful distinction between material characteristics. Materials with lumpy demand are characterised by both infrequent demand and high variability in demand quantities. This combination makes demand difficult to forecast and increases the risk of stockouts. As a result, safety stock levels may significantly improve material availability while reducing disruptions caused by unexpected demand peaks. Extracting the Class A materials from this category, creates a material group which is promising to test with.

3.4. Simulation-based optimisation framework

The proposed framework follows an **Optimisation with Simulation-based Iterations** (OSI) structure, shown in Figure 4. For a given set of safety stock parameters, multiple simulations are run, and each simulation consists of a sequence of weekly MRP-Opt executions representing the rolling-horizon planning process. Demand uncertainty is added by perturbing production plans using the estimated historical delta distributions. The final-week obsolete stock cost is used as the simulation outcome, and the average over all simulations is used as the optimisation objective value for a candidate parameter setting. To make safety stock a meaningful external optimisation parameter, the study introduces a **hard safety stock** (HSS) constraint. In the original model, safety stock is represented as a soft constraint with slack, which means the model may violate the intended stock level.

The proposed hard constraint removes that slack and enforces a true lower bound on inventory. Safety stock is then expressed in **days of sales** (DOS), which allows stock levels to be scaled consistently across materials with different demand characteristics. Candidate values are sampled within predefined bounds.

A **Full Enumeration** method is proposed as the approach for exploring the safety stock parameter space. This method evaluates every possible combination of parameter values within the predefined search domain. This search domain is bounded by 0 and 7 days of sales and it is explored for every integer value. By systematically testing all candidate solutions, it guarantees that the global optimal solution is identified, provided that the search space is finite and fully accessible. However, the computational cost grows rapidly with the number of optimisation parameters. Full Enumeration is therefore mostly practical for problems with a relatively small number of parameters or limited parameter ranges.

4. EXPERIMENTAL RESULTS

Several experiments are performed to calibrate the framework. First, the number of simulated weeks per run and the number of simulations per iteration are tested to balance reliability and runtime. Second, different variants of HSS and uncertainty assignment are explored. After this, varying artificial perturbation sizes are used in the production plans. A final proof-of-concept experiment aims to find a specific scenario in which HSS can be beneficial.

4.1. HSS and uncertainty variants experiments

To evaluate the proposed framework, several combinations of hard safety stock (HSS) assignment policies and uncertainty representations were tested. Four uncertainty settings were considered: no uncertainty, realistic uncertainty derived from historical production-plan variability as described in Section 3.2, deterministic perturbations applied during changeover weeks, and deterministic perturbations applied in the maximum changeover week. For each uncertainty setting, three HSS assignment variants were evaluated: constant HSS throughout the planning horizon, HSS active only during changeover weeks, and HSS active only in the maximum changeover week.

The results consistently showed that the lowest tested HSS level yielded the best performance among the HSS alternatives. However, none of the HSS variants outperformed the corresponding baseline without HSS, as can be seen in Table 1. Under realistic uncertainty, the baseline achieved an average obsolete stock cost of €70,044, whereas the best-performing HSS configuration resulted in €79,467 to €92,328 depending on the HSS assignment variant. Similar behaviour was observed under deterministic perturbation scenarios. Increasing HSS levels monotonically increased obsolete stock costs because additional inventory frequently remained unused after the changeover and became obsolete. These findings indicate that, when performance is measured solely through obsolete stock costs in the current conditions, imposing hard safety stocks does not improve system-wide performance.

4.2. Perturbation size experiments

To further investigate the interaction between uncertainty and safety stock, experiments were performed with deterministic demand perturbations of increasing magnitude. The perturbations were applied during the maximum changeover week, creating a

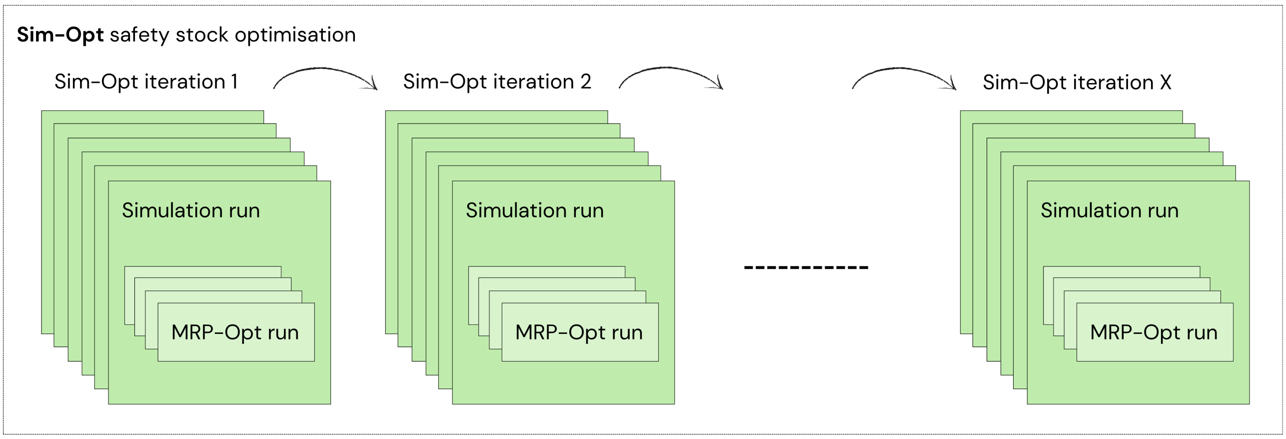

Figure 4. Sim-Opt framework

Table 1. Experiment results for different uncertainty and HSS configurations.

		Uncertainty			
		A	**B**	**C**	**D**
		No uncertainty	Realistic uncertainty based on PDFs	Perturbation in next simulation week when it is a changeover week	Perturbation in next simulation week when it is the max changeover week
HSS	No HSS	3 180	70 044	2 920	2 889
	Constant HSS	28 612	92 328	28 240	28 209
	HSS in changeover weeks	24 321	88 946	23 948	23 918
	HSS in max changeover week	24 905	79 467	24 471	24 502

stress scenario in which safety stock was expected to be most beneficial. Five representative items were selected, shown in Table 2, and perturbation sizes ranging from 123 to 123,456 units were tested.

The results showed that larger demand perturbations did not increase obsolete stock costs in the absence of HSS. On the contrary, sufficiently large perturbations reduced obsolete inventory because additional demand consumed stock that would otherwise become obsolete after the changeover. For the largest perturbations, obsolete stock costs even dropped to zero. When HSS was enforced, obsolete stock costs increased substantially for all perturbation sizes. Although larger perturbations consumed part of the buffer and slightly reduced the negative effect of HSS, they never offset the additional obsolete inventory created by the safety stock itself. Consequently, the no-HSS policy consistently outperformed all tested HSS levels.

Table 2. Perturbation size Experiments test items

Item [–]	**Changeover weeks** [–]	**Piece cost** [€]	**MOQ** [pieces]	**1 day of sales** [pieces]
1	2025-47; 2025-49	0.085176	5 760	48 703
2	2025-46; 2025-46	0.11044	5 446	9 876
3	2025-47; 2025-51	0.0064976	370 000	58 178
4	2025-46; 2025-52	0.06519	24 000	10 084
5	2025-52; 2025-52	0.2470703	2 700	4 004

4.3. Proof-of-concept experiment

A final proof-of-concept experiment was conducted to identify a specific scenario in which HSS could provide a benefit. The experiment focused on a single changeover item and combined targeted perturbations in the maximum changeover week with more finely grained HSS levels ranging from 0.1 to 7.0 days of sales.

The experiment confirmed the observations from the previous analyses. Large positive demand perturbations reduced obsolete stock by consuming more inventory before the changeover took place. However, once HSS became binding, the enforced buffer itself frequently remained after the changeover and was converted into obsolete inventory. Small HSS values below 0.3 days of sales had limited impact because they remained below the baseline stock level already present in the system. Larger HSS values increasingly dominated the final obsolete stock costs. Across all perturbation scenarios, the optimal solution remained the no-HSS policy.

These results demonstrate that the framework is capable of evaluating safety stock policies under uncertainty. However, they also show that additional safety stock does not provide measurable benefits under an obsolete-stock-cost objective. The findings suggest that any potential robustness advantages of safety stock are not captured by the current objective function and may require broader performance metrics that include inventory and emergency replenishment costs for example.

5. DISCUSSION

Although the proposed framework successfully integrated simulation-based optimisation with the existing MRP environment, the experiments did not demonstrate that hard safety stocks improve system performance. The most important explanation is the choice of objective function. Since performance was measured exclusively through obsolete stock costs, the framework captured the costs of additional inventory while largely ignoring potential benefits related to robustness, such as avoiding emergency replenishments or operational disruptions. As a result, hard safety stock was primarily penalised rather than rewarded.

The results also highlight the fundamental challenge of managing changeover materials. Unlike regular materials, changeover materials have a limited consumption window. Inventory that remains after a changeover immediately becomes obsolete, making additional stock inherently risky. The experiments showed that safety stock often remained unused and therefore increased obsolescence, even in scenarios designed to favour safety stock through artificial demand increases.

Another important observation concerns the aggregation of materials into groups. Grouping materials by demand pattern and economic importance made the optimisation problem computationally tractable, but it also reduced item-level specificity. Materials within the same group can differ substantially in cost, MOQ characteristics and demand behaviour. Consequently, a common safety stock policy may not reflect the requirements of individual materials.

Finally, this study illustrates the distinction between methodological feasibility and practical value. The framework proved capable of evaluating safety stock policies in a stochastic rolling-horizon environment and treating the industrial MRP model as a black box. However, the ability to optimise a parameter does not necessarily imply that the resulting policy creates measurable benefits under the selected performance metric.

6. CONCLUSION

This study investigated how simulation-based optimisation can be used to parametrise safety stocks of changeover materials in a deterministic MIP model operating in a stochastic MRP environment. To address this problem, a simulation-based optimisation framework was developed in which the existing industrial MRP optimisation tool was embedded within an Optimisation with Simulation-based Iterations structure. Demand uncertainty was modelled using historical production-plan data, and safety stock levels were optimised through a hard safety stock formulation.

The results show that simulation-based optimisation is a viable approach for evaluating and optimising planning parameters in a rolling-horizon MRP environment. The framework produced stable and reproducible results and successfully represented uncertainty while treating the MRP model as a black box.

However, no generic hard safety stock policy was found that improved performance. Across all experiments, increasing safety stock generally led to higher obsolete stock costs, while the best-performing solutions consistently corresponded to the lowest tested safety stock levels. Under the current objective structure, additional inventory introduced more obsolescence than it prevented.

The main contribution of this research is therefore the development and validation of a simulation-based optimisation framework for industry-scale MRP systems. Rather than identifying an optimal safety stock policy, the study clarifies the limitations of optimising robustness through obsolete stock costs alone and provides a foundation for future research into more comprehensive robustness measures.

7. FUTURE RESEARCH

Several opportunities exist to further develop the proposed framework. First, future research should broaden the objective function beyond obsolete stock costs. In the current study, safety stock is primarily penalised through additional obsolescence, while potential benefits such as preventing emergency replenishments, reducing rush-order costs, or improving service levels are not captured. Including these cost components would enable a more balanced evaluation of robustness.

Second, future studies should investigate more material-specific safety stock policies. The results suggest that a generic policy is unlikely to be effective across all changeover materials, as materials differ considerably in demand behaviour, economic importance, and replenishment characteristics. More targeted policies may better account for these differences and uncover cases where safety stock does add value.

Third, the uncertainty modelling could be further refined by using longer historical datasets and by incorporating seasonal effects, campaign-driven demand patterns, and other sources of variability. A more detailed representation of uncertainty may improve the realism of the simulation environment and provide a stronger basis for policy evaluation.

In addition, alternative robustness measures should be explored. Approaches such as safety lead times or MOQ-aware replenishment strategies may provide protection against uncertainty while reducing the risk of creating obsolete inventory.

Finally, future work should focus on improving computational scalability and validating the framework in practical planning settings. More advanced optimisation methods could enable the exploration of larger parameter spaces, while testing the framework on historical planning cases and comparing its recommendations with planner decisions would help assess its practical value.

AI Statement: During this study, Artificial Intelligence (AI) tools were used as supportive aids in several parts of the research and writing process. AI was used to assist with the development and refinement of Python code for modelling and experimentation, to support parts of the data analysis, and to help structure and report the results. In addition, AI was used to improve the readability of this paper and strengthen the use of scientific and formal academic language.

8. REFERENCES

[1] A. Dolgui and C. Prodhon, "Supply planning under uncertainties in MRP environments: A state of the art," *Annual Reviews in Control*, vol. 31, no. 2, pp. 269–279, 2007, ISSN: 13675788. DOI: 10.1016/j.arcontrol.2007.02.007.

[2] S. M. Javadi, S. J. Sadjadi, E. Teimoury, and A. Makui, "Material requirements planning with a novel lot sizing method and a new algorithm for production scheduling," *Scientific Reports*, vol. 15, no. 1, Dec. 2025, ISSN: 20452322. DOI: 10.1038/s41598-025-13197-8.

[3] H. Jodlbauer and A. Huber, "Service-level performance of MRP, kanban, CONWIP and DBR due to parameter stability and environmental robustness," *International Journal of Production Research*, vol. 46, no. 8, pp. 2179–2195, Apr. 2008, ISSN: 00207543. DOI: 10.1080/00207540600609297.

[4] R. D. Tordecilla, A. A. Juan, J. R. Montoya-Torres, C. L. Quintero-Araujo, and J. Panadero, *Simulation-optimization methods for designing and assessing resilient supply chain networks under uncertainty scenarios: A review*, Jan. 2021. DOI: 10.1016/j.simpat.2020.102166.

[5] J. Barros, P. Cortez, and M. S. Carvalho, *A systematic literature review about dimensioning safety stock under uncertainties and risks in the procurement process*, Jan. 2021. DOI: 10.1016/j.orp.2021.100192.

[6] R. Angkiriwang, I. N. Pujawan, and B. Santosa, "Managing uncertainty through supply chain flexibility: reactive vs. proactive approaches," *Production and Manufacturing Research*, vol. 2, no. 1, pp. 50–70, Jan. 2014, ISSN: 21693277. DOI: 10.1080/21693277.2014.882804.

[7] Z. Hong, C. K. Lee, and L. Zhang, *Procurement risk management under uncertainty: a review*, Sep. 2018. DOI: 10.1108/IMDS-10-2017-0469.

[8] W. Seiringer, K. Altendorfer, T. Felberbauer, B. Bokor, and F. Brockmann, "How periodic forecast updates influence MRP planning parameters: A simulation study," *Simulation Modelling Practice and Theory*, vol. 142, Jul. 2025, ISSN: 1569190X. DOI: 10.1016/j.simpat.2025.103115.

[9] S. Thevenin, Y. Adulyasak, and J. F. Cordeau, "Material Requirements Planning Under Demand Uncertainty Using Stochastic Optimization," *Production and Operations Management*, vol. 30, no. 2, pp. 475–493, Feb. 2021, ISSN: 19375956. DOI: 10.1111/poms.13277.

[10] M. Gansterer, C. Almeder, and R. F. Hartl, "Simulation-based optimization methods for setting production planning parameters," *International Journal of Production Economics*, vol. 151, pp. 206–213, 2014, ISSN: 09255273. DOI: 10.1016/j.ijpe.2013.10.016.

[11] B. Werth, J. Karder, A. Beham, and K. Altendorfer, "Simulation-based Optimization of Material Requirements Planning Parameters," in *Procedia Computer Science*, vol. 217, Elsevier B.V., 2022, pp. 1117–1126. DOI: 10.1016/j.procs.2022.12.310.

[12] W. Seiringer, J. Castaneda, K. Altendorfer, J. Panadero, and A. A. Juan, "Applying Simheuristics to Minimize Overall Costs of an MRP Planned Production System," *Algorithms*, vol. 15, no. 2, Feb. 2022, ISSN: 19994893. DOI: 10.3390/a15020040.

[13] W. Seiringer, K. Altendorfer, J. Castaneda, J. Panadero, and A. A. Juan, "Applying Simheuristics for Safety Stock and Planned Lead Time Optimization in a Rolling Horizon MRP System under Uncertainty," in *Proceedings - Winter Simulation Conference*, vol. 2021-December, Institute of Electrical and Electronics Engineers Inc., 2021, ISBN: 9781665433112. DOI: 10.1109/WSC52266.2021.9715294.

[14] G. Figueira and B. Almada-Lobo, "Hybrid Simulation-Optimization Methods: A Taxonomy and Discussion," Tech. Rep.

[15] D. B. Fontes, S. M. Homayouni, and J. F. Gonçalves, "A hybrid particle swarm optimization and simulated annealing algorithm for the job shop scheduling problem with transport resources," *European Journal of Operational Research*, vol. 306, no. 3, pp. 1140–1157, May 2023, ISSN: 03772217. DOI: 10.1016/j.ejor.2022.09.006.

[16] W. S. Lim and C. S. Tang, "Optimal product rollover strategies," *European Journal of Operational Research*, vol. 174, no. 2, pp. 905–922, Oct. 2006, ISSN: 03772217. DOI: 10.1016/j.ejor.2005.04.031.

[17] A. A. Syntetos, J. E. Boylan, and J. D. Croston, "On the categorization of demand patterns," *Journal of the Operational Research Society*, vol. 56, no. 5, pp. 495–503, 2005, ISSN: 01605682. DOI: 10.1057/palgrave.jors.2601841.

[18] A. A. Qaffas, M.-A. B. Hajkacem, C.-E. Ben Ncir, and O Nasraoui, "Interpretable Multi-criteria ABC analysis based on semi-supervised clustering and Explainable Artificial Intelligence," DOI: 10.1109/ACCESS.2022.0092316.