Journal Pre-proofs

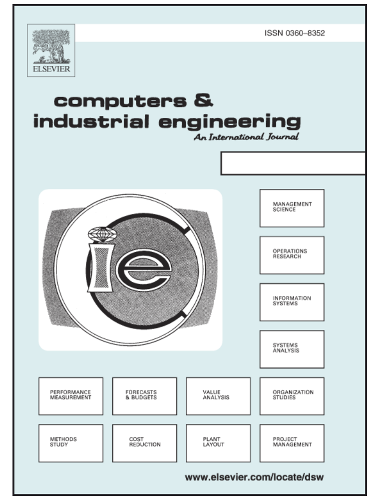

A hybrid bi-objective optimization approach for joint determination of safety stock and safety time buffers in multi-item single-stage industrial supply chains

Pedro M. Silva, João N.C. Gonçalves, Tiago M. Martins, Miguel Oliveira, Marcelo I. Reis, Luís Araújo, Daniela Correia, José Telhada, Lino Costa, João M. Fernandes

PII: S0360-8352(22)00165-6
DOI: https://doi.org/10.1016/j.cie.2022.108095
Reference: CAIE 108095

To appear in: *Computers & Industrial Engineering*

Received Date: 23 September 2021
Accepted Date: 9 March 2022

Please cite this article as: Silva, P.M., Gonçalves, J.N.C., Martins, T.M., Oliveira, M., Reis, M.I., Araújo, L., Correia, D., Telhada, J., Costa, L., Fernandes, J.M., A hybrid bi-objective optimization approach for joint determination of safety stock and safety time buffers in multi-item single-stage industrial supply chains, *Computers & Industrial Engineering* (2022), doi: https://doi.org/10.1016/j.cie.2022.108095

A hybrid bi-objective optimization approach for joint determination of safety stock and safety time buffers in multi-item single-stage industrial supply chains

Pedro M. Silva[1,*], João N.C. Gonçalves[1,2], Tiago M. Martins[1], Luís C. Marques[1], Miguel Oliveira[1], Marcelo I. Reis[1], Luís Araújo[1], Daniela Correia[1], José Telhada[2], Lino Costa[2], João M. Fernandes[3]

[1]Robert Bosch GmbH, Automotive Electronics Division, Divisional Logistics, Braga, Portugal
[2]ALGORITMI Centre, Department of Production and Systems, University of Minho, Braga, Portugal
[3]ALGORITMI Centre, Department of Informatics, University of Minho, Braga, Portugal

*Corresponding author: Pedro M. Silva[1] (pedrocmsilva97@gmail.com)

Acknowledgments

We would like to thank to the two anonymous reviewers for their helpful comments and suggestions. This work has been supported by the European Structural and Investment Funds in the FEDER component, through the Operational Competitiveness and Internationalization Program (COMPETE 2020) [Project no. 39479, Funding reference: POCI-01-0247-FEDER-39479].

"A hybrid bi-objective optimization approach for joint determination of safety stock and safety time buffers in multi-item single-stage industrial supply chains"

Highlights

September 23, 2021

- Safety stock and safety time inventory buffers are jointly optimized.
- Inventory holding costs are minimized while maximizing service level.
- A decision support system is implemented in a major automotive electronics company.
- The practical applicability of our approach is demonstrated on real-world data.

Portugal, September 23, 2021

Pedro M. Silva (Bosch Automotive Electronics, Portugal)

João N.C. Gonçalves (University of Minho / Bosch Automotive Electronics, Portugal)

Tiago M. Martins (Bosch Automotive Electronics, Portugal)

Luís C. Marques (Bosch Automotive Electronics, Portugal)

Miguel Oliveira (Bosch Automotive Electronics, Portugal)

Marcelo I. Reis (Bosch Automotive Electronics, Portugal)

Luís Araújo (Bosch Automotive Electronics, Portugal)

Daniela Correia (Bosch Automotive Electronics, Portugal)

José Telhada (University of Minho, Portugal)

Lino Costa (University of Minho, Portugal)

João M. Fernandes (University of Minho, Portugal)

A hybrid bi-objective optimization approach for joint determination of safety stock and safety time buffers in multi-item single-stage industrial supply chains

Abstract

In material requirements planning (MRP) systems, safety stock and safety time are two well-known inventory buffering strategies to protect against supply and demand uncertainties. While the role of safety stocks in coping with uncertainty is well studied, safety time has received only scarce attention in the supply chain management literature. Particularly, most previous operations research models have typically considered the use of such inventory buffers in a separate fashion, but not together. Here, we propose a decision support system (DSS) to address the problem of integrating optimal safety stock and safety time decisions at the component level, in multi-supplier multi-item single-stage industrial supply chains under dynamic demands and stochastic lead times. The DSS is based on a hybrid bi-objective optimization approach that simultaneously optimizes upstream inventory holding costs and β-service levels, suggesting multiple non-dominated Pareto-optimal solutions to decision-makers. We further explore a weighted closed-form analytical expression to select a single Pareto-optimal point from a set of non-dominated solutions, thereby enhancing the practical application of the proposed DSS. We describe the implementation of our approach in a major automotive electronics company operating under a myriad of components with dynamic demand, uncertain supply and requirements plans with different degrees of sparsity. We show the potential of our approach to improve β-service levels while minimizing inventory-related costs. The results suggest that, in certain cases, it appears to be more cost-effective to combine safety stock with safety time compared to considering each inventory buffer independently.

Keywords: Safety stock, Safety time, Multi-objective optimization, Decision support.

1. Introduction

A Supply Chain (SC) is generally understood as a network system that interconnects multiple players (e.g., suppliers, manufacturers, retailers) aiming to provide products, services and information that bring added value to the customers while minimizing related costs (Simchi-Levi et al., 2007; Lambert, 2008; Barbosa-Póvoa et al., 2018). In such a setting, the success of all SC players naturally depends on how well they interact with each other and account for business changes in volatile markets. The management of such interactions is commonly referred to as supply chain management (SCM) (Lambert and Cooper, 2000). Many of the efforts to increase customer satisfaction occur in the context of demand and inventory management (Fisher, 1997; You and Grossmann, 2008). In the latter context, manufacturing companies usually operate under Material Requirements Planning/Enterprise Resource Planning (MRP/ERP) systems for inventory replenishment (Whybark and Williams, 1976; Axsäter, 2006; Louly and Dolgui, 2013). To cope with demand and supply uncertainties inherent to such systems, companies often take advantage of safety stock and/or safety time inventory buffers. While the first translates into an additional inventory beyond the regular stock, the second consists in planning order releases earlier than established in the requirements plan and scheduling their receipt earlier than required (Whybark and Williams, 1976; Alves et al., 2004). Determining the proper safety buffering level for each product is held to be one of the most robust strategies to soften uncertainty (Koh et al., 2002) and has been investigated extensively given its importance to the operations research (OR)/management science community in general.

From an optimization perspective, it is well-known that optimal inventory buffers should be suitably computed according to the trade-off between inventory-related costs and service level requirements. In this context, previous research has shown important theoretical guidelines for choosing between safety stock and safety time in MRP systems characterized by demand and supply uncertainty (either in terms of quantity or timing) (Whybark and Williams, 1976; Grasso and Taylor III, 1984; Lambrecht et al., 1984; Chang, 1985; Etienne, 1987; Buzacott and Shanthikumar, 1994; Van Kampen et al., 2010). However, while recent years have seen the increasing development of OR modeling approaches for dimensioning safety stocks in many inventory control contexts (Gonçalves et al., 2020), the literature is scarce in respect to strategies for optimal parameterization of safety time – which, in industrial practice, is often set based on experience (Louly et al., 2008; Dickersbach, 2009). More importantly, by typically considering the use of safety stocks and safety times in a separate fashion, the literature has provided limited insight into how these two buffering strategies might be used in combination for achieving target service levels at minimum cost. In sharp contrast, we are interested in examining whether and in what circumstances the combination of both buffering strategies could be a more cost-effective approach than considering either safety stocks or safety times in isolation. Our interest also lies in providing insights into how this combination depends on the magnitude of demand variation and supply delay, as well as the degree of sparsity of material requirements plans.

To address these shortcomings, we develop a hybrid bi-objective optimization model to jointly optimize safety stock and safety time decisions in multi-supplier multi-component single-stage automotive electronics supply chains with dynamic demands and uncertain supply. We turn our attention to the manufacturer stage of the SC (hereinafter referred to as company), which is naturally subject to upstream and downstream variabilities. We take inventory holding costs and β-service level as primary objectives to be optimized. The safety time and safety stock solutions generated during the optimization process are successively evaluated, in terms of the proposed objectives, via a MRP simulation model. Afterwards, we explore a closed-form analytical expression that enables to select a single Pareto-optimal point from a comprehensive set of non-dominated solutions, by weighting different key logistics performance indicators. Throughout this work, our formulations relax the common assumptions that demand and supply lead time are Gaussian-distributed random variables, which may not always be realistic and may represent an oversimplification of the real SC dynamics. Motivated by the lack of large-scale empirical studies in this context, we developed and implemented a decision support system (DSS) comprising the proposed bi-objective optimization approach for optimal safety buffer parameterization.

Our contributions to the area of inventory management are fourfold:

1. We propose a DSS, comprising a hybrid bi-objective optimization approach, to jointly determine safety stock and safety time buffers in MRP systems.

2. We investigate how both the sparsity of delivery schedules and the supply/demand variability impact on the choice of the optimal cost-effective inventory buffer.

3. We provide guidelines to decision-makers and practitioners interested in the optimal parameterization of safety inventory buffering strategies – a fundamental problem affecting the effectiveness of MRP (Den Boer, 1994; Louly and Dolgui, 2013).

4. We demonstrate the practical applicability and the financial/operational benefits of

the proposed DSS by considering a real-world case study from a major multi-item multi-supplier automotive electronics company.

In the rest of this paper, we start by motivating our work in this context by providing some application examples from the literature (Section 2). Then, we introduce the mathematical optimization model to jointly optimize safety stock and safety time in Section 3. Section 4 presents and describes the system architecture of the DSS that supports the proposed mathematical model and facilitates its use in real-world operational contexts. The proposed approach is tested with illustrative examples from the case study company and the results derived therefrom are discussed in Section 5. Finally, we highlight the practical and theoretical implications of our work and summarize its conclusions in Section 6.

2. Related work

In general MRP production systems, supply and demand processes are the two basic sources of variability, commonly expressed in the form of timing and/or quantity (Whybark and Williams, 1976). In such systems, safety stock and safety time buffering methods are the primary means to cope with variability factors (Guide Jr and Srivastava, 2000). In what follows, we provide a general yet non-exhaustive literature overview on existing modeling approaches to estimate safety stocks and safety times. We focus on the problem of dimensioning safety buffers for each component. Readers interested in other variants of the problem, like safety stock placement or positioning, are pointed to the fundamental works of Inderfurth and Minner (1998); Minner (2001, 2012); Caridi and Cigolini (2002); Graves and Willems (2000, 2003); Graves and Schoenmeyr (2016) and references cited therein.

2.1. On the safety stock and safety time estimation problem

In real-world inventory control topologies, demand is not deterministic and it should be forecasted (Prak et al., 2017). As such, for a given target service level, classical inventory management theory typically establishes safety stocks based on the standard deviation of past demand forecasting errors over a lead time. Recent published research (Saoud et al., 2021; Babai et al., 2022) provides interesting guidance on the estimation of demand over lead time, given its role for dimensioning safety stocks. In this context, theoretical and parametric/nonparametric empirical approaches are suitable strategies to estimate the variation of demand forecast errors (see, Silver et al., 2016; Trapero et al., 2019a,b, for more details).

Since the fundamental study of Eppen and Martin (1988), a number of studies have been undertaken to address the safety stock dimensioning problem from different perspectives, namely considering (Persona et al., 2007; Gonçalves et al., 2020): the standard deviation of normally distributed demands (Inderfurth and Vogelgesang, 2013; Lu et al., 2016; Braglia et al., 2016; Bahroun and Belgacem, 2019), the standard deviation of demand forecasting errors (Trapero et al., 2019a,b; Kanet et al., 2010; Beutel and Minner, 2012; Boute et al., 2014), the role of product structure and component standardization (Collier, 1982; Carlson and Yano, 1986; Molinder, 1997; Grubbström, 1998) or even the introduction of additional uncertainty issues beyond the classical ones related to demand and supply (Cheung and Hausman, 1997; Chakraborty and Giri, 2012; Sarkar et al., 2010). A substantial body of literature also assumes independent and identically distributed (iid) Gaussian forecast errors when dimensioning safety stocks. However, in practice, if this assumption is violated, the service levels may be subject to positive and negative fluctuations, leading to an increase in inventory-related costs (Beutel and Minner, 2012; Trapero et al., 2019a,b). We refer the interested reader to the recent studies of Trapero et al. (2019a,b) for more details on how to bypass such classical assumptions.

Sharing the same goal as the safety stocks, safety time inventory buffers are characterized by bringing forward orders and their respective receipts earlier than agreed upon in the MRP requirements plan (Whybark and Williams, 1976). On the choice between safety stocks and safety times, several studies have been providing guidelines to choose the most adequate buffering method for improving customer service level while reducing inventory holding costs. There is, however, no general agreement as to the method that should generally be used (see Guide Jr and Srivastava, 2000, for a comprehensive review and discussion on this topic). Whybark and Williams (1976) found that, regardless of the source of variability, safety stock is suitable for quantity uncertainty whereas safety time is preferred for timing uncertainty. Yet, while Sato and Tsai (2004) favor the use of safety time buffers as a more flexible technique than safety stocks in most situations where timing uncertainty prevails, Grasso and Taylor III (1984) suggest the use of safety stocks to tackle uncertain supply/timing. It is noteworthy that, by focusing on supply and demand variability in isolation, the above works do not evaluate the joint use of both sources of variability. In sharp contrast, Etienne (1987) proposed a model for choosing the most cost-effective buffering method considering the existence of uncertainty in both supply and demand sides. It is shown that safety time should be disregarded for MRP production systems operating under quantity variability. Moreover, concerning timing variability, safety time proves to be useful in sparse schedules in detriment of safety stock. The work of Etienne (1987) further underlines, yet not surprisingly, that there is no such buffering method with superior performance in every situation. Molinder (1997) also considers both supply and demand uncertainty factors to explore, via simulated annealing, how the uncertainty level affects the optimal choice of the buffering strategy. The main conclusion is that safety time works best whenever both supply and demand are highly variable – a result further corroborated by Van Kampen et al. (2010) – while safety stocks are recommended for cases with high demand variability and low lead time variability. Interestingly, a previous work of Buzacott and Shanthikumar (1994) argued that safety time buffers are preferred in detriment of safety stocks whenever demand is properly forecasted, favoring the use of safety stocks when demand variability increases. This disagreement may arise from different modeling assumptions considered. While the model of Buzacott and Shanthikumar (1994) does not include supply variability in their model formulations, Molinder (1997) considers both supply and demand variability factors. More recently, Avci and Selim (2016, 2017, 2018) developed simulation-based models that account for safety time to solve the inventory replenishment problem in different supply chain topologies. Simulation-based procedures have also proven to be useful to effectively design virtual safety stocks under given inventory control systems (Dickersbach, 2009; Nenni and Schiraldi, 2013). Other interesting recent studies (Ben-Ammar et al., 2019, 2018) address the problem of optimal parameterization of planned lead times which, together with forecasted lead time, can be used to derive better safety time estimations.

Despite the lack of consensus, the previous literature has suggested relevant insights that may help grasp how and in what situations buffering methods are used. Unfortunately, the overwhelming majority of the literature does not rely on real-world data from industrial case studies, an issue that had already been raised in the fundamental work of Guide Jr and Srivastava (2000).

2.2. Highlights from the literature and motivation

We found that empirical evidence on the joint optimization of safety stock and safety time is rather scarce, and the combined used of both buffers has proved to be effective in practice (Glock and Kim, 2016; Kania et al., 2021). Motivated by the scarce industrial case-study applications in this context, especially with multi-component multi-supplier considerations, it is our belief that there is a significant potential in the joint exploration of safety time and safety stock buffering mechanisms.

At a first sight, the joint use of both buffering strategies may seem redundant. Yet, note, for instance, that a theoretical increase in one-day of safety time for a component with a once-per-week receiving frequency might force a company to push the delivery order backwards by one week, so that to ensure that the planned order reception match a planning calendar day contracted with the supplier (see Fig. 1). Consequently, this naturally implies a huge increase in inventory holding costs, by forcing the company to bring forward the scheduled order delivery well ahead of time. On the other hand, the effectiveness of safety time is also dependent on the supplier delivery performance. In cases where the supply delay is greater than the established safety time, this safety margin may be insufficient to cover the supply variability. Here, instead of increasing safety time, it would be interesting to evaluate the potential of maintaining (or decreasing) it and introduce suitable quantities of safety stock. We hypothesize that this combination may prove to be suitable for balancing holding costs while maintaining target service levels, especially for manufacturing components with very sparse receiving frequencies. Alternatively, one could obviously relax the use of safety time and take fully advantage of safety stock to cope with demand and supply variabilities. Nevertheless, the holding costs derived from this strategy may be very high and, depending on the degree of uncertainty at the upstream and downstream end of the SC, some part of this safety stock may never be used. All of these arguments justify the interest of combining both safety inventory buffers in a multi-objective optimization fashion.

In a nutshell, this research intends to better understand the (operational/financial) benefits arising out of the joint optimization of safety stock and safety time buffers. For that, we propose a hybrid bi-objective optimization model embedded in a DSS to propose optimal buffering strategies for components with different supply, demand and MRP dynamics, providing support to logistics managers and planners in their planning daily operations.

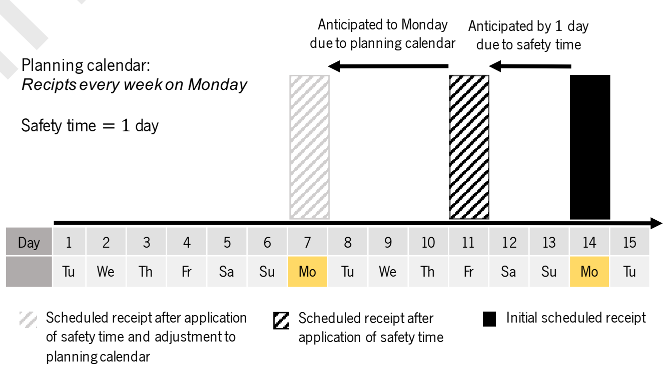

Figure 1: The effect of increasing safety time margins in sparse material requirement plans.

3. Problem description and modeling

3.1. Preliminary concepts

Before elaborating on the proposed mathematical optimization model, let us provide the essential background on multi-objective optimization that is relevant throughout the paper. The familiar reader with this topic is invited to skip to Section 3.2.

A general multi-objective optimization problem can be generally defined as (Deb, 2005):

$$\begin{aligned} \min \quad & f(\mathbf{x}) = (f_1(\mathbf{x}), \ldots, f_p(\mathbf{x})) \\ \text{s.t.} \quad & \mathbf{x} \in S \end{aligned} \tag{1}$$

where $S \subseteq \mathbb{R}^n$ is the decision (or feasible) set, $\mathbb{R}^n$ is the decision space, and $f_i : S \to \mathbb{R}$, for each $i = 1, \ldots, p$, are the objective functions to be minimized. One can extend the formulation (1 to cases where some or all objective functions are to be maximized. The feasible set can be characterized by nonlinear inequalities, equalities and bounded constraints, i.e., $S = \{\mathbf{x} \in \mathbb{R}^n : g_r(\mathbf{x}) \geq 0, \forall r = 1, \ldots, R,\ c_k(\mathbf{x}) = 0, \forall k = 1, \ldots, K,\ \underline{\mathbf{x}}_l \leq \mathbf{x}_l \leq \overline{\mathbf{x}}_l, \forall l = 1, \ldots, L\}$. Given $S' \subseteq S$, we denote by $\mathbf{f}(S') = \{\mathbf{z} \in \mathbb{R}^p : \exists \mathbf{x} \in S'$ such that $\mathbf{z} = \mathbf{f}(\mathbf{x})\}$ the image set of S' in the objective (or solution) space $\mathbb{R}^p$, with $\mathbf{f}(\mathbf{x}) = (f_1(\mathbf{x}), \ldots, f_p(\mathbf{x}))$. Likewise, $Z = \mathbf{f}(S)$ is the image set of S in the solution space $\mathbb{R}^p$, which consists of all feasible solutions (or points) derived from the multi-objective optimization problem (1).

The conflicting nature of the objectives implies that it is, in practice, impossible to optimize them simultaneously. Given two feasible points $\mathbf{z}, \mathbf{z}' \in Z$, we say that $\mathbf{z}$ dominates $\mathbf{z}'$ ($\mathbf{z} \prec \mathbf{z}'$) iff

$$\mathbf{z}_i \leq \mathbf{z}'_i,\ \forall i \in \{1, \ldots, p\} \wedge \exists j \in \{1, \ldots, p\} : \mathbf{z}_j < \mathbf{z}'_j\ . \tag{2}$$

Similarly, we say that $\mathbf{z}$ weakly dominates $\mathbf{z}'$ ($\mathbf{z} \preceq \mathbf{z}'$) iff $\mathbf{z}_i \leq \mathbf{z}'_i,\ \forall i \in \{1, \ldots, p\}$. We denote by $N = \{\mathbf{z} \in Z : \nexists\ \mathbf{z}' \in Z$ with $\mathbf{z}' \prec \mathbf{z}\}$ the set of non-dominated solutions (also called *Pareto set* (Deb, 2005)). In this context, one can also define the *ideal* ($\mathbf{z}^*$) and *nadir* ($\mathbf{z}^{nad}$) points to represent the best and worst objective values, respectively:

$$\mathbf{z}^* = \min_{\mathbf{z} \in Z} \{\mathbf{z}_i\}_{i \in \{1, \ldots, p\}}\ ,\ \mathbf{z}^{nad} = \max_{\mathbf{z} \in Z} \{\mathbf{z}_i\}_{i \in \{1, \ldots, p\}}\ . \tag{3}$$

3.2. General description and assumptions

The SC topology underlying the inventory problem of this study consists of a single company, operating with multiple suppliers and components, following a MRP methodology for inventory replenishment. We consider that the MRP system operates under a rolling frozen period over the planning horizon, as shown in Fig. 2, in which changes in the production schedule during the frozen period are generally not allowed. In such a setting, we are interested in jointly optimizing, at each iteration of the rolling scheme, safety stock and safety time decisions for each component c in such a way that upstream inventory holding costs are minimized and the β-service level to production is maximized. To be consistent with the reasoning behind the MRP methodology, optimal decision scenarios are provided only for the time horizon immediately starting from the end of the frozen period (denoted by t_c^f) onwards. In short, safety optimal solutions for each component are proposed for all time periods in the set $\mathcal{T} = \{t : t_c^f \leq t \leq T_c\}$, where T_c is the maximum optimization-simulation horizon.

Throughout the paper, we assume that each component c is supplied by a single supplier following a specific supply policy. Components under vendormanaged inventory (VMI) and consignment policies are out of scope. The company's demand for each component c is assumed to be dynamic over a finite planning horizon of discrete time periods, obtained from bill-of-materials (BOM) explosions on finished product forecasts provided by downstream SC players. Against this background, the company places orders to suppliers to fulfill end-customer requirements.

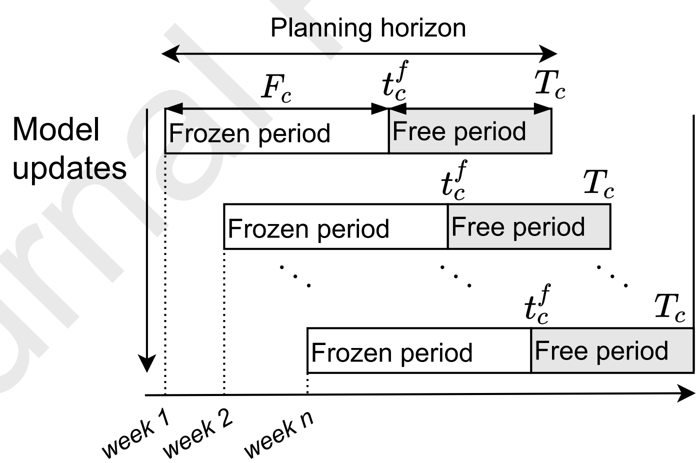

Figure 2: Rolling MRP planning horizon strategy with frozen and free periods for a component c.

Table 1 summarizes the notation used throughout this paper to denote sets, parameters, decision variables and functions.

Table 1: Summary of notations used throughout the paper.

Sets	
$\mathcal{C} = \{1, \ldots, \vert\mathcal{C}\vert\}$	Set of components
$\mathcal{T}_0 = \{1, \ldots, \vert\mathcal{T}_0\vert\}$	Set of periods preceding the end of the frozen period
$\mathcal{T} = \{1, \ldots, \vert\mathcal{T}\vert\}$	Set of periods following the end of the frozen period
$\mathcal{P} = \{1, \ldots, \vert\mathcal{P}\vert\}$	Set of periods matching order receiving dates
Parameters	
t_c^f	Period matching the end of the frozen period for component $c \in \mathcal{C}$
T_c	Maximum simulation horizon ($T_c \in \mathcal{T}$) for component $c \in \mathcal{C}$ (days)
F_c	Length of frozen period for component $c \in \mathcal{C}$ (days)
$D_{c,t}$	Company's demand for component $c \in \mathcal{C}$ in period t (units)
M_c	Minimum order quantity for component $c \in \mathcal{C}$ (units)
$O_{c,t}$	Scheduled order quantity receipt for component $c \in \mathcal{C}$ in period t (units)
X	Supplier delay risk; a discrete random variable (days)
$I_{c,t}$	Inventory level for component $c \in \mathcal{C}$ in period t (units)
h_c	Inventory holding cost of a component $c \in \mathcal{C}$ (Euro per unit per period)
γ	Relative change factor of demand for component $c \in \mathcal{C}$
Decision variables	
ST_c	Safety time for component $c \in \mathcal{C}$ (days); an integer decision variable
SS_c	Safety stock for component $c \in \mathcal{C}$ (units); an integer decision variable
Functions	
$\mathbb{E}[\cdot]$	Expectation operator
$1_{\{A\}}$	Indicator function (equals 1 if A is true and 0 otherwise)
$(\cdot)^+$	Maximum function, $(\cdot)^+ = \max(\cdot, 0)$

3.3. Mathematical model formulation

We consider a bi-objective mathematical optimization model to jointly optimize safety time and safety stock buffers for each component $c \in \mathcal{C}$, in terms of the classical cost-service trade-off in inventory management theory. This model is formulated as follows:

$$\min_{ST_c, SS_c} \quad H(ST_c, SS_c) = \frac{h_c}{T_c - t_c^f + 1} \sum_{t=t_c^f}^{T_c} I_{c,t}(ST_c, SS_c) \tag{4}$$

$$\min_{ST_c, SS_c} \quad U(ST_c, SS_c) = \frac{\sum_{t=t_c^f}^{T_c} \left(D_{c,t} - I_{c,t}(ST_c, SS_c)\right)^+}{\sum_{t=t_c^f}^{T_c} D_{c,t}} \tag{5}$$

subject to:

$$\underline{ST}_c \leq ST_c \leq \overline{ST}_c, \ \forall c \in \mathcal{C} \tag{6}$$

$$\underline{SS}_c \leq SS_c \leq \overline{SS}_c, \ \forall c \in \mathcal{C} \tag{7}$$

$$I_{c,t}, D_{c,t} \geq 0, \ \forall c \in \mathcal{C}, \forall t \in \mathcal{T} \tag{8}$$

Objective (4) minimizes the total averaged inventory holding costs over the optimization-simulation period, whereas objective (5) minimizes the total average fraction of unmet company's demand, which is logically equivalent to maximizing average company's service level. Backlog costs are not considered in the formulation due to the difficulty to measure them in real-world contexts (Petropoulos et al., 2019). For this reason, we prefer the use of a β-service level approach, implicitly considered in objective function (5), to estimate the expected fraction of total company's demand that can be fulfilled from the inventory on-hand. Concretely, in objective (5), the fraction of fulfilled company's demand for c is maximal if $D_{c,t} \leq I_{c,t} \forall t \in [t_c^f, T_c]$. In contrast, a shortage of $D_{c,t} - I_{c,t}$ inventory units occurs whenever $D_{c,t} > I_{c,t}$. By measuring the magnitudes of the shortages during the optimization-simulation window $[t_c^f, T_c]$, one can thus derive the overall fraction of company's demand that cannot be fulfilled during that period. Here, particular attention should be given to the term $I_{c,t}(\cdot, \cdot)$, which is a function of the decision variables and thereby plays a fundamental role in both objectives. This function aims to estimate the inventory levels for component c over the optimization-simulation period if we vary the safety time and safety stock values. One can interpret this function as a proxy for the expected inventory dynamics derived from a MRP simulation in a given ERP system using the safety buffering parameterization (ST_c, SS_c). Section 4.2.2 provides the details on how this function is embedded in a simulation-based optimization procedure by taking into account dynamic demand and uncertain supply. Constraints (6) and (7) define the upper and lower bounds of the decision variables and (8) ensures the non-negativity of the state variables.

By solving the above-mentioned bi-objective optimization problem with suitable evolutionary metaheuristics (see Section 4.2.2), we aim at generating the trade-off curve between the two conflicting objectives (4) and (5), allowing decision-makers to choose a Pareto-optimal solution according to their preferences.

4. Proposed decision support system

We developed a DSS in order to facilitate the use of the previously proposed bi-objective optimization model in a real-world supply chain context. Such a DSS was conceived in the course of the project *Data Intelligence and Analytics for Business Operations* (DIABO), resulting from a partnership between Bosch AE and University of Minho, Portugal. Next, we follow the work of Deng et al. (2021) to structure and describe the system architecture of the proposed DSS.

The DSS herein presented builds on a three-layer architecture comprised by a database, a simulation-based optimization model and a graphical user interface (GUI) (Fig. 3). The database layer aggregates all the MRP-related information and data concerning the different types of components targeted for safety buffering optimization. Such information is extracted from a central database that acts as a copy of the company's ERP and serves as the data foundation for all modeling actions. This layer also comprises a database for storing the output derived from the simulation-based optimization process. The information extracted from the databases is thereafter processed in the process layer, the core component of the entire system architecture, which is responsible for jointly optimizing the safety stock and safety time levels for each component according to the objective functions defined in Section 3.3. Finally, the proposed DSS has a flexible GUI layer that allows the user to interact with the system in different stages.

Each of these layers is thoroughly explained in the subsequent sections.

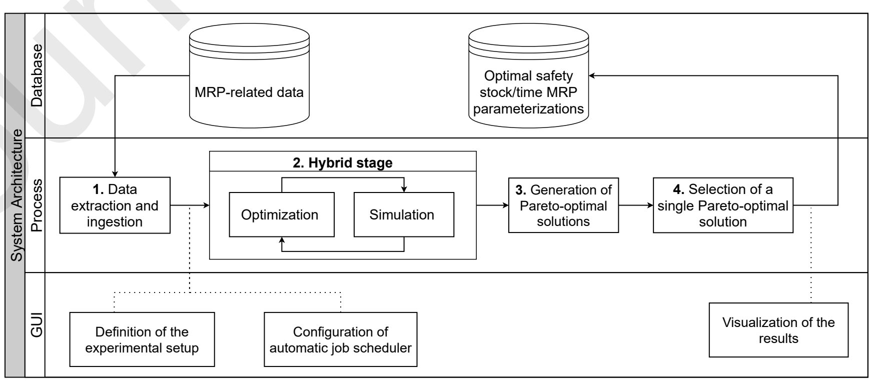

Figure 3: The system architecture of the proposed DSS.

4.1. Database layer

The database stores MRP-related data. Such data attributes fall into the standard inputs required to run a classical MRP system (Ptak and Smith, 2016), namely the bill-of material data (component's structure), inventory data, components/suppliers-related master data, as well as the master production schedule (MPS) information. The data attributes collected were further subject to a more detailed exploratory data analysis, by checking for missing values, outliers, consistency and completeness. The cleansed data resulting from this procedure serve as input for building the MRP simulation system described in Section 4.2.2, which is tasked to simulate the impact, in terms of upstream inventory holding costs and β-service levels, of changing safety stock and/or safety time as MRP parameters for each manufacturing component.

4.2. Process layer

The process layer comprises the simulation-based optimization procedure adopted to generate the set of non-dominated optimal safety solutions for the different components. This section starts by describing the extraction and ingestion processes of the data collected from the database layer previously presented (Section 4.2.1), which serve as data feeding mechanisms for running the bi-objective optimization model presented in Section 3. The optimization and simulation phases comprising such model are thereafter described in Section 4.2.2. We further introduce a weighted closed-form analytical expression to select a single Pareto-optimal solution from a set of Pareto-optimal points, according to different key performance indicators (Sections 4.2.3–4.2.4). Finally, we provide details on the GUI embedded in our DSS (Section 4.3).

4.2.1. Data extraction and ingestion

Although the architecture portrayed in Fig. 3 presents a single database, thereby simplifying the acquisition and subsequent integration of the information, a massive amount of MRP-related data associated with thousands of components can systematically be updated into the system. To cope with the need for faster processing of data with heterogeneous formats, we take advantage of a Hadoop cluster for Big Data processing[1]. In such a setting, the data are firstly extracted via Apache Sqoop (Ting and Cecho, 2013) and ingested

[1]The interested reader on the use of Big Data Analytics in the context of SCM is referred to the works of Wang et al. (2016); Zhong et al. (2016).

in a Hadoop Distributed File System (HDFS) (Shvachko et al., 2010). The Apache Spark (Zaharia et al., 2016) engine is then used for data processing, including for the distributed implementation of the hybrid bi-objective evolutionary-based optimization model described in Section 4.2.2. The enriched data obtained from the processing stage are further loaded and stored in Hive (Thusoo et al., 2009) tables, each corresponding to a HDFS directory. At this point, one can leverage this data repository to rapidly query, via Impala (Russell, 2014), the solutions of the optimization model and design data visualization tools to enhance the user experience and support decision-making.

4.2.2. Hybrid stage

In the context of inventory management, it is known that pure analytical/optimization models are, in general, difficult to implement in real-world SCs (Avci and Selim, 2018). As such, our DSS comprises a simulation-based optimization (hybrid) approach in order to better mimic the effect of re-parameterizing safety time and safety stock parameters in a MRP inventory replenishment system and further evaluate such parameterization in terms of the proposed conflicting optimization objectives. Our approach follows a two-stage modeling process where the generation of the set of non-dominated Pareto-optimal solutions is obtained through an iterative process between a selected algorithm providing feasible solution pairs (ST_c, SS_c) (optimization stage) and a deterministic simulation module that evaluates each solution in terms of the proposed objectives (simulation stage). Similar optimization-simulation setups are adopted in other multi-objective optimization applications (see Avci and Selim, 2017, 2018; Altazin et al., 2020). Following a classical evolutionary approach, the proposed hybrid approach appears summarized in Fig. 4, and the optimization and simulation stages comprising it are detailed below.

Step 1: Optimization stage

In order to compute the non-dominated solutions for bi-objective optimization problem, we consider three popular (Nebro et al., 2009a; Durillo et al., 2010) Pareto dominance based evolutionary computation metaheuristics, namely the Non-dominated Sorting Genetic Algorithm II (NSGA-II) (Deb et al., 2002), the Muti-Objective Cellular (MOCell) genetic algorithm (Nebro et al., 2009b) and the improved version of the Strength Pareto Evolutionary Algorithm (SPEA), called SPEA2 (Zitzler et al., 2001). A brief overview on each technique is provided as follows.

The NSGA-II is considered to be the one of the most popular population-based meta-

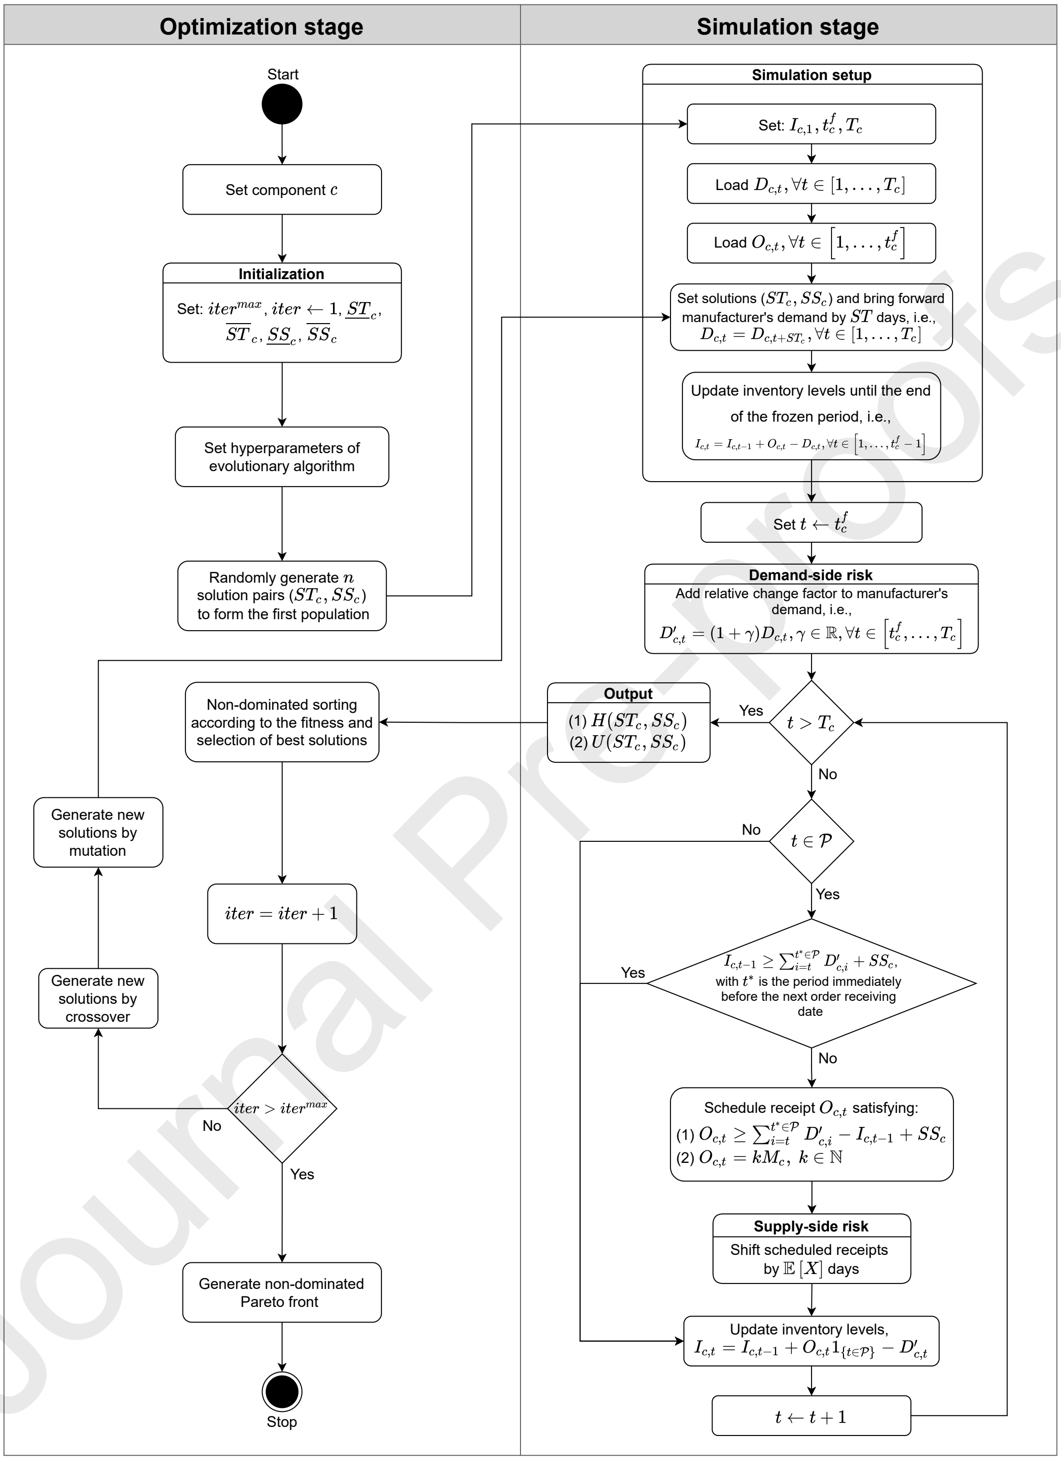

Figure 4: The proposed optimization-simulation approach for the joint optimization of safety time and safety stock buffers over a given planning period.

heuristics for multi-objective optimization (Li and Zhang, 2008), and takes advantage of a fast non-dominated sorting approach to rank the different solutions according to Pareto dominance and optimality concepts (Deb et al., 2002) introduced previously in Section 3.1. If two solutions share the same ranking (i.e., are non-dominated to each other), the one with the highest crowding distance is preferred in such a way that widespread and diversity of solutions are preserved.

The MOCell is a cellular genetic algorithm that operates under the reasoning that each individual only interacts with individuals from its neighborhood. In particular, it stores a set of non-dominated solutions in an *external archive* during the search procedure (Nebro et al., 2009b) using the same crowding distance of NSGA-II. By selecting two neighbor parent solutions, crossover and mutation operators are used to generate and a new individual. MOCell employs a feedback strategy from the archive to the population so that a new individual is replaced by an archived solution if the former is worse (in the Pareto sense) than the latter. Analogously to MOCell, the SPEA2 algorithm also includes an external archive to store non-dominated solutions that result from the application of genetic operators. Yet, the SPEA2 employs an improved truncation method whenever the number of non-dominated solutions is greater than the population size so that solutions with minimum distance to any other solution are preferred to those with higher distances (Nebro et al., 2009a; Zitzler et al., 2001).

Following the optimization stage depicted in Fig. 4, the initialization process includes the definition of several classical features in any evolutionary-based optimization metaheuristic, ranging from the solution encoding scheme and the fitness function to the algorithm input control parameters, including the (i) the population size (n) and the bounds of the decision variables, (ii) the maximum number of iterations, (iii) the selection mechanism, and (iv) the genetic operators (mutation and crossover) with respective probabilities (Turan et al., 2020). For all genetic algorithms described previously, we adopt the standard binary tournament (Deb et al., 1995) for the selection procedure, and the simulated binary crossover (SBX) and polynomial mutation (Deb, 2005) for the crossover and mutation genetic operators, respectively, with distribution indexes denoted by η_c and η_m. These indexes govern the spread of offspring solutions around parent solutions (Deb, 2000). The selection of suitable probability values for the genetic operators (p_m and p_c for the probability of mutation and crossover) is typically performed by following a tuning procedure (e.g., Taguchi approach (Roy, 2001)) or manual experimentation.

After the initialization process, all the steps presented in Fig. 4 form the generic workflow of a classical genetic algorithm, including: the generation of an initial set of feasible random solutions (forming the initial population), the solutions evaluation and subsequent application of genetic operators (selection, crossover and mutation). In this whole process, note that the simulation module is invoked at each iteration of the optimization process to evaluate and provide the objective values for each solution vector generated during the optimization stage. The above process operates in a cyclic fashion to successively generate Pareto frontiers until some stopping criteria is met (in Fig. 4, a predefined maximum number of iterations, $iter^{max}$).

Step 2: Simulation stage

As articulated in the previous step, each feasible solution pair (ST_c, SS_c) generated during the optimization stage is evaluated through a simulation procedure to estimate both the expected inventory holding cost and the fraction of unmet company's demand derived therefrom. In short, this simulation procedure intends to mimic the MRP logic embedded in any ERP system, in terms of inventory and replenishment dynamics, for a given safety time and safety stock parameterization.

Given the relevance of the simulation stage in the evaluation of feasible solutions generated during the optimization stage, we consider convenient to describe the steps outlined in the right-hand part of Fig. 4 by narrative, for an arbitrary component c.

We start by setting the fundamental setup parameters for the simulation procedure, including the inventory at the beginning of the planning horizon ($I_{c,1}$), the day matching the end of the frozen period (t_c^f) and the maximum simulation horizon (T_c). Next, the forecasted company's demand derived via BOM-explosions ($D_{c,t}$) is loaded to the simulator for all time periods up until T_c and brought forward in time by ST days, meaning that $D_{c,t} = D_{c,t+ST}, \forall t \leq T_c$. This latter step forces the MRP model to plan order receipts earlier, thereby matching the definition and overall purpose of the safety time buffer. By loading the scheduled receipts in the frozen period, together with $D_{c,t}$ and initial inventory $I_{c,1}$, it is possible to update the inventory levels from $t = 1$ up to $t = t_c^f$, inclusively. This step enables the computation of the initial inventory at the beginning of the free period, whereupon our optimal solutions are proposed according to our bi-objective optimization problem.

The second phase of the simulation is started immediately after the end of the frozen period. From this point onwards, supply and demand uncertainty are considered within the simulation scheme. Demand-side risk is firstly included by adding a (positive/negative) correction factor γ to the company's demand forecast after the frozen period. For that purpose, we compare the demand volume planned for the upcoming week after the end of the frozen period at two different moments: at the beginning of the frozen period ($t = t_c^f - F_c$) and at the end of the frozen period ($t = t_c^f$). This allows to compute the relative change in demand for a specific period. The final correction factor γ is thus obtained by aggregating past demand relative changes without outliers. Although not a guarantee against future demand uncertainty, this correction factor allows, at a certain level, to adjust upcoming recent demand forecasts. Next, for each valid order receiving date in the supplier planning calendar, we determine whether a supplier order receipt $O_{c,t}$ should be scheduled in that period. A supplier order is scheduled at time t if the available inventory-on-hand $I_{c,t-1}$ does not cover the sum of company's demand until the day immediately before the next valid delivery date t^* in the planning calendar plus the safety stock. If an order is scheduled, the respective quantity should be a multiple of the minimum order quantity ($O_{c,t} = kM_c, k \in \mathbb{N}$) while covering the total expected net demand the safety buffer SS_c.

To account for supply timing uncertainty related to a component c, we model the historical supplier delay risk by a random variable X_c described by a discrete probability distribution with a finite support consisting of different magnitudes of past order's delay, ranging from X_c^- to X_c^+, and respective probabilities. At this point, we assume that past recent supply dynamics hold into the future. Then, we consider that each scheduled receipt may be delayed by approximately $\mathbb{E}[X_c]$ days, meaning that an order scheduled to be delivered at period t may be actually received at period $t + \mathbb{E}[X_c]$, potentially leading to shortages depending on the magnitude of the delay and the inventory-on-hand. As a conservative approach, we consider that the greatest delay in the random variable X_c is applied to the order for which the inventory-on-hand is minimal, whereas the remaining scheduled order receipts suffer a positive time deviation with magnitude $\mathbb{E}[X_c]$. After incorporating the supplier delay risk in each scheduled receipt, the inventory records for the current period t are updated accordingly and serve as basis for developing the requirements plan in the following period $t + 1$.

The above processes are repeated iteratively over $t \in [t_c^f, T_c]$, and the objective values for each solution vector (ST_c, SS_c) under testing are then successively returned to the optimization stage.

4.2.3. Generation of Pareto-optimal solutions

For each component $c \in \mathcal{C}$, the simulation-based optimization procedure derives a set of non-dominated Pareto-optimal solutions. Each solution in the Pareto front corresponds to a particular Pareto-optimal safety time/stock solution, with a certain holding cost and β-service level associated. From this, the decision-maker disposes of a panoply of solutions to choose from, each one associated to distinct inventory trade-offs on the two evaluation criteria defined. For each Pareto front generated, we compute the hypervolume indicator (or size of space covered, in Zitzler and Thiele (1998, 1999)) to evaluate the performance of the genetic algorithms when generating the final population set across the experimental studies. We recall that in a bi-objective minimization problem, given a set $N = \{\mathbf{z}^a, \mathbf{z}^b, \ldots, \mathbf{z}^y\}$ of non-dominated solutions, the hypervolume consists of the measure of the objective space which is concomitantly dominated by N and bounded above by a reference point $r \in \mathbb{R}^2$ such that $r \geq \mathbf{z}^{nad} = \max_{\mathbf{z}\in N}\{\mathbf{z}_i\}_{i\in\{1,2\}}$, with the relation $\geq$ being applied in a componentwise fashion (Fonseca et al., 2006).

4.2.4. Selection of single of Pareto-optimal solutions

Complex SCs typically operate with multiple components of different nature from multiple suppliers worldwide, which makes it unfeasible for logistics planners to devote particular attention to each individual item. In such a scenario, choosing the optimal safety parameterization for each component from a set of Pareto-optimal solutions may prove to be a time-consuming task for the management.

For a component c, let $N \subseteq \mathbb{R}^2$ be the set of non-dominated Pareto-optimal safety time/safety stock pairs in the form $\mathbf{s} = (ST_c, SS_c)$. Inspired by the holistic measure of inventory performance proposed by Petropoulos et al. (2019), we have constructed a flexible weighted closed-from analytical expression aiming to simplify the choice of an optimal safety buffer MRP parameterization $\mathbf{s}^* \in N$ from a set of Pareto-optimal solutions. It should be emphasized that the process of choosing such an optimal parameterization $\mathbf{s} \in N$ for each component takes place only after the generation of the corresponding Pareto-optimal front. Hence, the choice of such an optimal parameterization has no impact on the optimization-simulation process. The proposed expression encompasses the linear combination of four performance indicators to be minimized and it is formulated as follows:

$$\mathbf{s}^* = \underset{\mathbf{s}\in N}{\arg\min} \sqrt{W_1\left(\frac{H(\mathbf{s})}{\overline{H}}\right)^2 + W_2\left(\frac{U(\mathbf{s})}{\overline{U}}\right)^2 + W_3\left(\frac{F(\mathbf{s})}{\overline{F}}\right)^2 + W_4\left(\frac{C(\mathbf{s})}{\overline{C}}\right)^2}, \quad (9)$$

where $W_i, i = 1, \dots, 4$ are weight factors for the different indicators. The first two terms of Eq. (9) are related to the evaluation criteria considered in the bi-objective optimization problem formulated in Section 3.3. Hence, we seek a Pareto-optimal solution $\mathbf{s}$ for which the corresponding holding costs and unfulfillment rates are minimized. We further explore two additional performance indicators apart from $H(\cdot)$ and $U(\cdot)$. The first is the expected average premium freight cost generated by the adoption of a certain Pareto-optimal pair ($F(\mathbf{s})$). Recalling the formulation of the objective function (5), the adoption of unsuitable safety buffer levels may lead to specific rates of unmet company's demand, which may compel the company to resort to premium transportation freights in order to request the amount of stock in shortage and satisfy the demand production requirements $D_{c,t}$, $\forall t \in [t_c^f, T_c]$. At this point, note that special freight costs may not evolve linearly with the inventory in shortage, as they depend on other factors related to size, weight and geographical location of potential suppliers. For this reason, we consider interesting to include this indicator as a complement to the indicator $U(\cdot)$, in the sense that a high magnitude of inventory in shortage may not necessarily lead to high special freight costs. Similarly to the indicators $H(\cdot)$ and $U(\cdot)$, we are interested in a Pareto-optimal solution $\mathbf{s}$ such that $F(\mathbf{s})$ is also minimized. The second additional indicator targeted for minimization is the inventory coverage (in time) provided by a certain Pareto-optimal pair ($C(\mathbf{s})$), here formulated as the difference between the inventory coverage (in days) provided by the safety solution $\mathbf{s}$ and the expected value of supplier delay $\mathbb{E}[X_c]$. Following such formulation, we seek a solution $\mathbf{s}$ that provides enough inventory time coverage to cope with the average delay without however incurring in excessive holding costs. Here, only Pareto-optimal points satisfying $C(\mathbf{s}) \geq \mathbb{E}[X_c]$ are potential candidates to be considered as $\mathbf{s}^*$. Note that, depending on the context of application, other business rules can be applied in order to reduce the set of potential candidate solutions to be selected as $\mathbf{s}^*$. Overall, while the indicator $F(\cdot)$ complements $U(\cdot)$, the indicator $C(\cdot)$ may help to balance the minimization process of holding costs ($H(\cdot)$) to a reasonable level enough to cover supply timing uncertainty.

As the different criteria in formulation (9) are measured in different units (costs, percentages and time), each criterion (term) is normalized by considering its value divided by the mean value of that criterion over all Pareto-optimal solutions $\mathbf{s} \in N$. Our formulation considers the L_2-norm in order to cope with diseconomies of scale in operations management (see Petropoulos et al., 2019, and the references cited therein for details). The choice of the weight parameters in Eq. (9) can be configured by the user according to their business preferences through a graphical user interface layer (see Section 4.3.1). Particularly, the user can opt to allocate more weight to one specific indicator than another in accordance with the optimization objective that has greater relevance at a given moment in time. If the user intends to favor the minimization of inventory holding costs in detriment of the minimization of the fraction of unmet demand (or, equivalently, the maximization of the service level), then the magnitudes of W_1 and W_4 should be higher than those assigned to the factors W_2 and W_3. Such configuration of weights will promotes the selection of lower safety buffers. Conversely, if the purpose is to focus on minimizing the fraction of unmet demand regardless of the inventory holding costs derived therefrom, then the magnitudes of W_2 and W_3 should be higher than those assigned to W_1 and W_4, thereby promoting the selection of solutions with higher safety coverage.

4.3. Graphical user interface layer

The GUI layer allows the user to interact with the DSS. Following the functional requirements of the DIABO project, we distinguish three functionalities embedded in our DSS that allow the user to: (i) define the experimental setup for the simulation-based optimization model, (ii) configure an automatic job scheduler and (iii) visualize the Pareto-optimal safety stock/time solutions for each component, as well as some related performance indicators.

4.3.1. Definition of the experimental setup

The DSS presents a main interface comprising different functionalities of the system (Fig. 5). To match the purpose of this study, we focus our attention only to the safety time/stock optimization and configuration functionalities, indicated by the circle markers 1 and 2 in Fig. 5, respectively. The latter functionality is now introduced, whereas the former is detailed later in Section 4.3.3.

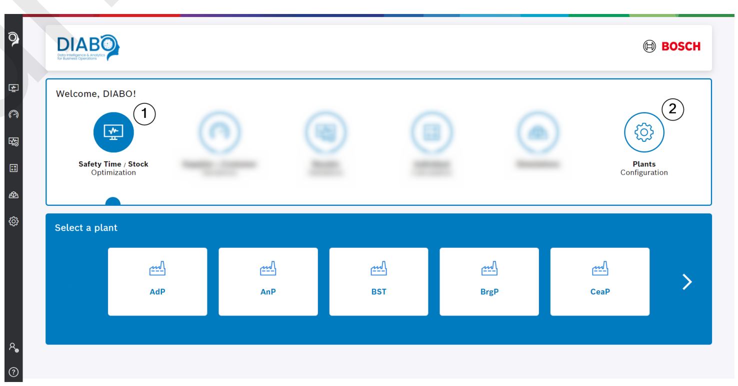

Figure 5: Main interface of the DSS.

The "Plants Configuration" button, indicated by the circle marker 2 in Fig. 5 enables to configure a number of settings that are relevant for running the hybrid bi-objective optimization model. By clicking on it, the DSS displays a list of different plants that can be targeted for such configurations. Figure 6 provides an illustrative example of a truncated list of plants. The user can then select the edit button marked by 3 in Fig. 6 in order to dive into the configuration page for a specified plant. Several configurations are available, for instance: (i) define a small subset of components of that plant to be optimized as well as the corresponding optimization-simulation horizon, (ii) configure automatic alert mechanisms to logistics planners in case of abnormal supply/demand deviations, (iii), set specific safety time/stock conservative lower bounds per component (typically motivated by business constraints), or even (iv) define the hyperparameters of the metaheuristic and the weights for the closed-form analytical expression (9). Such configurations can be conducted by accessing the sub-panel represented by the circle marker 4 in Fig. 6.

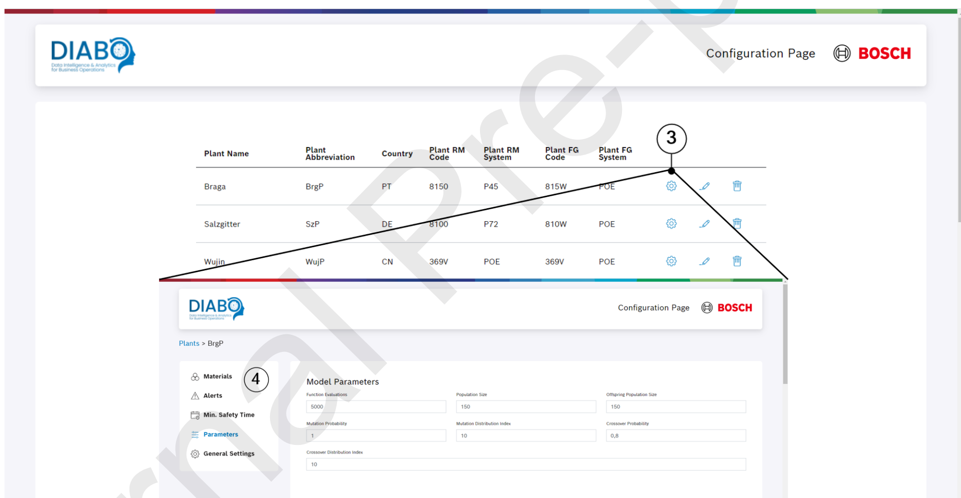

Figure 6: Configuration page of the DSS.

4.3.2. Configuration of the automatic job scheduler

The GUI allows the automatic scheduling of several jobs required for running the execution of routines associated with the bi-objective optimization model for different plants and components. Such scheduling process can be accessed directly through the side-bar menu button indicated by marker 5 in Fig. 7, which displays several examples of jobs scheduled for different time periods. The user is also able to schedule a new job by clicking the button indicated by the marker 6 in Fig. 7, as well as to monitor the status progress of each scheduled job and have a direct visual assessment of the processes whose execution may have failed. In practice, this feature is important as it enables to quickly detect possible errors in order to further conduct corrective actions.

Prior to launch a given job in a productive environment (marker 8 in Fig. 7) the user is able to run it in a test environment (marker 7 in Fig. 7) in order to evaluate its performance and validate the results derived therefrom.

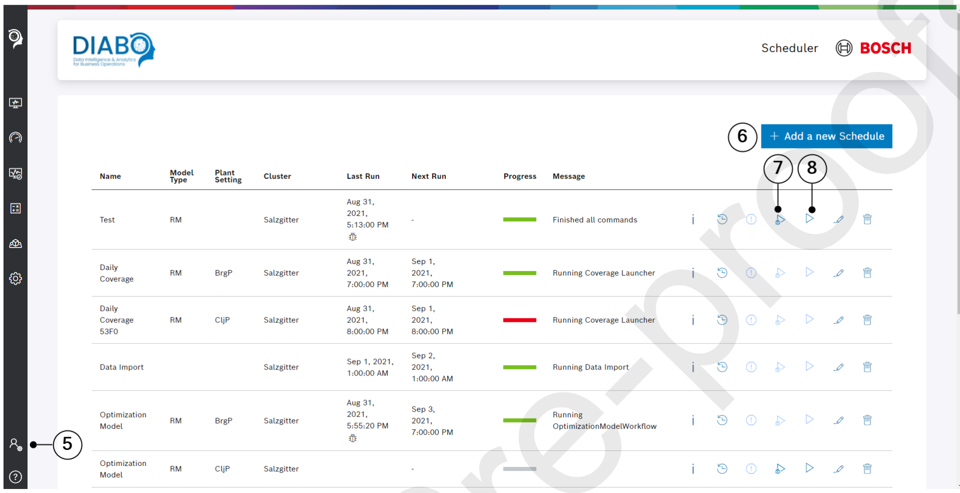

Figure 7: Interface of the automatic job scheduler.

4.3.3. Visualization of the results

The results of the bi-objective hybrid optimization model are presented to the user through the button indicated by marker 1 in Fig. 5. The output is illustrated in Fig. 8. In addition to the visualization of the Pareto-optimal front for a selected component (see marker 9), the user can select a particular Pareto solution and analyze the inventory dynamics obtained by adopting the corresponding buffering strategy over the simulation period (see marker 10). In practice, this enables logistics planners to visually inspect the inventory management impact of a given safety buffer choice, thus making it possible to anticipate potential inventory shortfalls and consequent service level losses. The Pareto-optimal solution selected via the closed-form expression (9) is represented by the green point in the Pareto front, as indicated in the panel corresponding to marker 9. Since the component presented in Fig. 8 accounts for a significant share of the company's sales volumes, it was suggested by the company managers to give greater importance to the expected service level than to the related holding costs. Thus, following the reasoning articulated in Section 4.2.4, the weights in Eq. 9 were set, for this particular component, as $W_1 = W_4 = 0.1, W_2 = W_3 = 0.4$.

This leads to the selection of a Pareto-optimal point associated with higher safety coverage and, consequently, higher service level (as shown in Fig. 8). In the panel corresponding to marker 9, one can also visualize the current safety time/stock solution adopted by the company (yellow point in the sub-panel 9), for the sake of comparison with the Pareto-optimal solutions. Several master data related to the component under observation are likewise presented to the user, including the variations in supply and demand variations associated with that specific component, as indicated by marker 11.

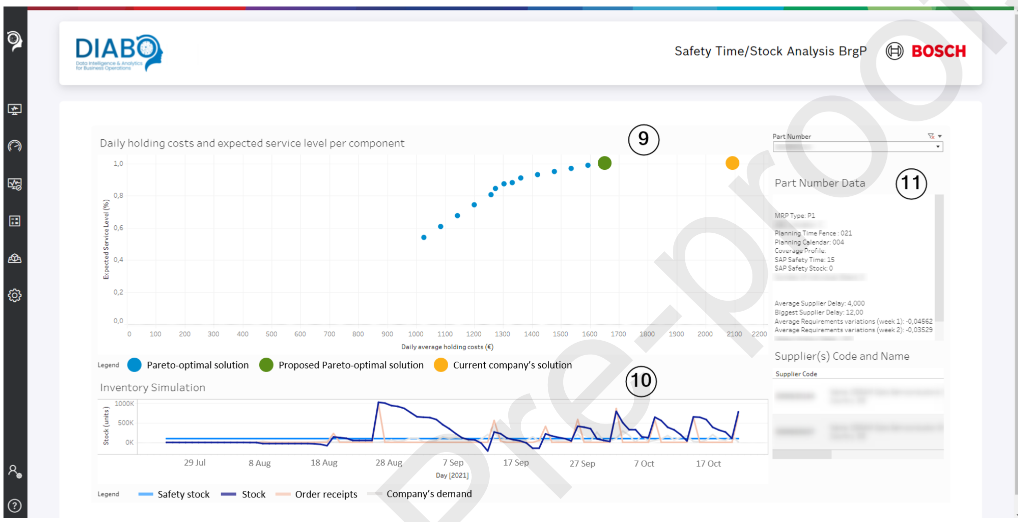

Figure 8: Output of the safety stock/time optimization functionality.

5. Empirical evaluation

To evaluate the merits of our approach, the DSS was implemented in the logistics department of Bosch AE, Portugal – one of the largest business units belonging to the Bosch group, operating with several suppliers and end-customers worldwide. Presently, the company dimensions safety buffers according to business-oriented rules closely related to the ABC inventory classification of the different components and to the geographical location of potential suppliers. Yet, the company managers are interested in understanding how the dimensioning of the safety buffers is influenced by uncertainty factors as well as by the characteristics of component's supply from suppliers to the manufacturing plant. Based on this background, we have designed two test cases. The first to perceive how the joint optimal parameterization of safety buffers is affected not only by the magnitude of demand and supply variations, but also by the degree of sparsity of the corresponding MRP plans.

Regarding the second test case, we aimed to compare, at a more generic level, the financial and operational benefits resulting from the combined used of safety stock and safety time with those obtained by the current approach adopted by the company. In both test cases, to cope with the current company's strategy, we guide the discussion by relating the results derived from the proposed simulation-based optimization approach with the ABC classification of the different manufacturing components under consideration. At this point, we consider a sample of 3024 components of a single manufacturing plant.

5.1. Calibration of models and parameters

All the metaheuristics for solving the proposed bi-objective optimization problem (NSGA-II, MOCell, and SPEA2) have been implemented in Python by taking advantage of the flexible jMetalPy package (Benitez-Hidalgo et al., 2019). The computation experiments were performed on a Hadoop cluster with 9 nodes, 152 cores (144 cores with hyper-threading), 224 TB of disk capacity and 1603 GB of memory capacity.

We adopted an internal population size of 150 individuals, for all the evolutionary metaheuristic employed. We set the distribution indexes for the crossover and mutation genetic operators to $\eta_c = \eta_m = 10$, and the size of the archive of solutions to 100, in the case of MOCell and SPEA2. Throughout the numerical experiments, the crossover probability is assumed to be $p_c = 0.8$ and the mutation probability is $p_m = 1/I$, where $I = 2$ is the number of decision variables at stake. For all metaheuristics, we set the maximum number of iterations to $iter^{max} = 1500$, as we found through background analyses that the average hypervolume is quasi-invariant from 1500 iterations onwards. However, in order to save some computational time, we adopt a dynamic stopping criterion rather than forcing the model to iterate until the maximum number of iterations is exhausted. Specifically, every five generations of the model, the number of solutions in the current Pareto front is evaluated and compared with the number of solutions in the front obtained five generations ago. Thus, the iterative process stops whenever the percentage change in the number of solutions from one Pareto front to another is less or equal than a threshold value (in our case 2%). The iterative process continues until this criterion is satisfied or the maximum number of iterations set ($iter^{max}$) is reached.

Given the fairly large dimension of our component's sample, grid search procedures for the hyperparameters of the genetic algorithms were not conducted to enable a better computational efficiency. Yet, we found through preliminary experiments that small perturbations in the hyperparameters produced quite identical simulation results. Regarding the lower and upper bounds of the decision variables, we consider that ST_c is a discrete variable taking values between $\underline{ST}_c = 0$ and $\overline{ST}_c = \lceil \max(m, X_c) + \mu_{D_c}/\sigma_{D_c} \rceil$ with a unit step size, where m is the maximum safety time defined at the company and μ_{D_c}/σ_{D_c} is the coefficient of variation of the past annual company's demand. On the other hand, SS_c ranges from a lower bound $\underline{SS}_c = 0$ up to an upper bound $\overline{SS}_c = \max\{O_{c,t}\}_{t=1}^{t_c^f}$, with a step size of M_c. The setup for these bounds was established based on expert business judgment for this particular case-study context. As such, there may be naturally the need to redefine such bounds depending on the context under analysis. The maximum simulation horizon for each component c is set to $T_c = 10$ (weeks). We considered one-year historical data to model the demand-side variations and the supply delay risk, defined as in the simulation stage (see Section 4.2.2). For the computation of the hypervolume metric, the objective space was normalized so that the nadir and ideal points are $(1, 1)$ and $(0, 0)$, respectively, taking the former as reference point.

5.2. Results and discussion

5.2.1. Preliminary experiments

A set of controlled experimental studies were performed to select the top performing evolutionary algorithm for our data, thereby reducing the computational effort required to optimize all the manufacturing components contained in our sample using the three different metaheuristics. For that, we took advantage of a set of 30 components with high turnover rate for the company. We have considered both the hypervolume and runtime (in seconds) as model performance metrics. The three metaheuristics (NSGA-II, MOCell and SPEA2) were applied to each component over 5 runs, with 1500 function evaluations in each one of them. The results were then aggregated by averaging the results across runs. A final estimated median for the hypervolume and runtime for the whole set of components was obtained via non-parametric Wilcoxon signed-rank test (Hollander et al., 2013), as summarized in Table 2.

Table 2: Wilcoxon medians for the hypervolume and runtime for the different evolutionary algorithms.

Metaheuristic	Hypervolume	Runtime (s)
NSGA-II	0.665	13.483
MOCell	0.664	191.367
SPEA2	0.666	262.748

Although no statistical significant differences ($p > 0.05$) were found for the hypervolume values among the different algorithms used, we found that the computational time is considerably smaller for NSGA-II when compared to the MOCell and SPEA2. For this reason, the NSGA-II was selected as the fastest metaheuristic for the data under study. Subsequent analyses consider the application of NSGA-II with the hyperparameters defined in Section 5.1.

5.2.2. Influence of schedule density on the Pareto-optimal decision space

We start by studying the influence of schedule density on the dynamics of the Pareto-optimal pairs (ST, SS) obtained from the application of NSGA-II. Henceforth, we define schedule density as the relative frequency of scheduled receipts for a given component over the optimization-simulation period. In other words, a manufacturing component with a dense planning calendar can be characterized by frequent scheduled deliveries over the simulation horizon, while another with a sparse calendar has scheduled deliveries in wider timeframes. In what follows, we consider only Pareto-optimal solutions providing service levels $\geq 90\%$. This threshold of 90% for the service level was established in accordance with the minimum service level which the case-study company is willing to comply with, and agrees with previous literature (Fildes and Kingsman, 2011). We distinguished the component types according to their ABC inventory class. Following this strategy, we expect that our analyses are not overly influenced by components with very different inventory management dynamics. The three inventory groups (A, B and C) were obtained based on the annual consumption value criterion.

In Fig. 9, we illustrate the dynamics of the Pareto-optimal solutions over the decision space as a function of schedule density and ABC class. Due to scaling reasons, safety stock is expressed in days of inventory coverage rather than units. Each point in the decision space represents a given Pareto-optimal solution for a particular component in the original sample. For components within the classes A and B (left and middle panels of Fig. 9), we observe the overwhelming majority of the solutions combining safety stock and safety time decisions are mainly concentrated in decision space coordinates corresponding to lower degrees of schedule density. This suggests that, in terms of the proposed optimization trade-off, the combination of safety time and safety stock decisions may be suitable for components for which supplier's deliveries are more widely separated in time, i.e., with sparse planning schedules. Note that the adoption of a pure safety time strategy for components with sparse planning schedules would force the company to anticipate supplier orders well-ahead in time, inducing highly holding cost levels. Instead, this opens space to reduce the levels of safety time and compensate them with suitable safety stocks.

As we move from sparse schedules to dense schedules (i.e., high degree of density), we found that the need to combine both buffering strategies seems to lose steam. Indeed, one can argue that, in the absence of unusual demand/supply variability, it is reasonable to soften safety levels whenever a given component is supplied more frequently. For very dense schedules (density ≥ 0.9), we found that the need for safety buffering strategies is residual, regardless the ABC class which the components belong to. A closer comparison between the dynamics of the Pareto-optimal solutions within class C with those of the remaining classes shows that this typology of components tends to require less safety time. Such behavior can be motivated by the close distance between the supplier of those components and the manufacturing plant under study. In this regard, attending to the low inventory holding cost of C-components, logistics planners can thus dimension suitable safety stock buffers, thereby reducing the risk of experiencing inventory shortages for commodities and bulky materials resulting from ordering well beforehand.

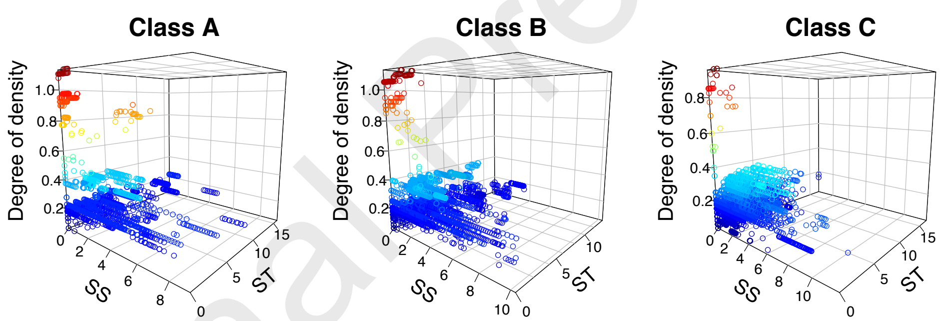

Figure 9: Decision space dynamics as a function of the degree of density.

5.2.3. Influence of uncertainty sources on the Pareto-optimal decision space

Let us focus now on the potential link between the proposed optimal solutions and the different levels of uncertainty, either in the demand or in the supply side. Likewise, we narrow our analyses to service levels greater or equal to 90%.

The influence of demand variation on the dynamics of the decision space, under different ABC classifications, is depicted in Fig. 10. On one hand, by examining the general interplay of the solutions obtained for the different ABC classes, it is clear that the percentage demand variation tends to be negative for the overwhelming majority of the components tested. This means that, for our data, demand signals tend to be more often overestimated than underestimated. We found that such behavior was motivated by the conservative behavior adopted by the logistics planners of Bosch AE in a bid to cope with the major supply chain disruptions induced by SARS-CoV-2 virus at the time of this study. On the other hand, we observed that components with highly negative demand variations levels tend to present Pareto-optimal solutions with low magnitudes of safety stock, and only a few of those solutions is complemented with safety time. In sharp contrast, whenever the demand variation positively increases, components tend to require the combined usage of safety stock and safety time to reduce the potential risk of stock-out events. This behavior is particularly evident for components of class A (left panel of Fig. 10).

As to the decision space dynamics as a function of supply uncertainty (Fig. 11), here expressed in the form of timing uncertainty, we found no strong patterns that would allow us to extract meaningful conclusions. Yet, one can observe that components associated with high supply delays which fall into classes A and B tend to be associated with Pareto-optimal solutions that combine safety stock with safety time decisions. This pattern is not so evident for C-components.

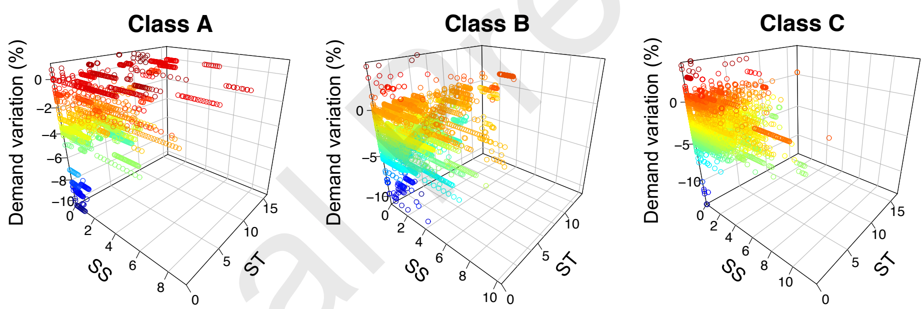

Figure 10: Decision space dynamics as a function of demand quantity uncertainty.

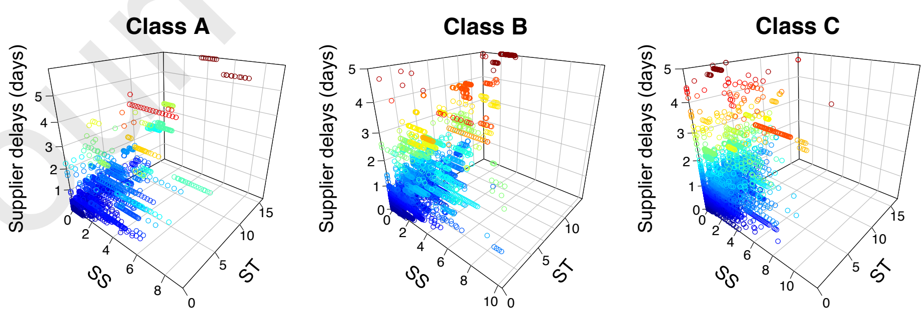

Figure 11: Decision space dynamics as a function of supply timing uncertainty.

5.2.4. On the benefits of the joint optimization of safety inventory buffers

Currently, safety stock and safety time inventory-buffering strategies are adopted by the company. However, company managers consider the dimensioning of both buffers in isolation and neglect the potential benefits arising out of their joint optimization. On the other hand, such dimensioning process is purely empirical, stemming from the experience of logistical planners, and based on ABC criteria and the location of suppliers. Hence, we are now interested in quantifying the benefits resulting from the joint optimization of safety inventory buffers in comparison to that obtained from the current strategy adopted by the case-study company.

We started by analyzing the distributions of the solutions obtained through the proposed optimization-simulation approach with those obtained by using the company's strategy (Fig. 12). In general, we observe that the average inventory coverage provided by the combination of safety time and safety stock buffers is lower when compared to that obtained by using the current company's strategy, regardless the ABC class. We further observe that a significant part of the inventory coverage levels are obtained by using safety time compensated with safety stock. In this context, it additionally important to evaluate the financial and operational performance of the proposed approach. For that, we compared, for all sampled components, the current safety solution adopted by the company with the Pareto-optimal point obtained via the application of the formulation (9), assuming equal weights $W_i = 0.5, \forall i = 1, \ldots, 4$. The results appear summarized in Table 3. The key finding is that whenever safety buffers are jointly optimized the resulting holding costs are always lower and the average service levels are always higher, when compared to those obtained by considering safety buffers dimensioned in a independent fashion. At first glance, this result may seem a contradiction because an increase in the service level for a given component is typically associated with an increase in inventory levels. However, by considering the joint optimization of both buffers, we observed that, for multiple components, it was possible to drastically reduce inventory costs while maintaining the same service level derived from the strategy adopted by the case company. This leads to another interesting finding: although there are several Pareto-optimal decisions associated with increases in safety levels in relation to the current ones adopted by the company, we found that the benefits arising from such cost reduction outweigh the additional holding costs induced by increasing safety stock/time. Particularly, the greatest improvement in service levels was found for components in class A, whereas in terms of holding costs the largest reduction was found for components in class C. This last finding is mainly due to the fact that the company tends to naively increase safety stock/time levels for components within class C, as they present considerably lower holding costs when compared to those of classes A and B. While this strategy adopted by the case-study company may be pertinent from an operational point of view, it may not be cost-effective if the increase in safety inventory levels is too excessive.

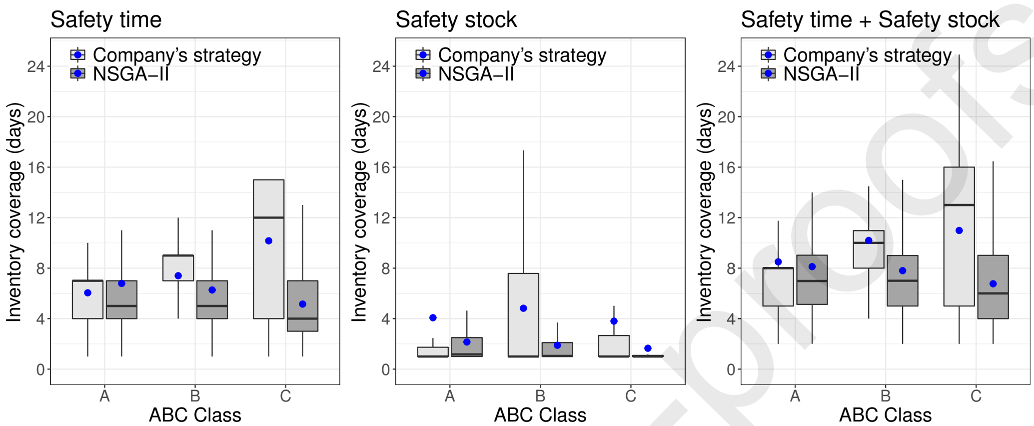

Figure 12: Inventory coverage distribution of all Pareto-optimal safety time (ST) and/or safety stock (SS) solutions derived from the simulation-optimization procedure via NSGA-II (dark gray) and the case company's safety strategy (light gray) for the different ABC classes. The blue dots represent average values.

Table 3: Comparison of average holding costs and service levels derived from the joint buffer optimization with those obtained under the case company's strategy for different ABC classes.

	Performance indicator	Class A	Class B	Class C
NSGA-II	Average β-SL	0.85	0.95	0.99
	Average holding cost [€]	10 200	795	86
Company's strategy	Average β-SL	0.75	0.94	0.97
	Average holding cost [€]	11 300	954	114
	Overall β-SL improvement [%]	13.33	1.06	2.06
	Overall holding cost savings [%]	9.73	16.67	24.56

6. Conclusions and further developments

The literature has emphasized that no supply chain can operate without safety inventory buffers (Syntetos et al., 2016), as they are required at different levels of the product structure in order to guarantee timely production and delivery of finished products and components

(Louly and Dolgui, 2009; Ruiz-Torres and Mahmoodi, 2010), as well as to prevent massive supply chain disruptions (Bogataj et al., 2016). Yet, several previous studies have typically looked at the problem of dimensioning safety stocks or safety times in isolation and not in combination, thus ignoring any potential inherent to the combination of these two safety approaches.

From a theoretical perspective, we introduce a hybrid bi-objective optimization model in order to address the problem of joint optimization of safety stock and safety time inventory buffers in multi-item single stage industrial supply chains. Classical evolutionary metaheuristics are employed to generate multiple Pareto-optimal safety stock/time solutions to decision-makers. Here, instead of giving the decision-maker the onus of choosing the Pareto-optimal solution that works best for the business needs, we explored a holistic metric that, by weighting different performance indicators, allows to suggest a single Pareto-optimal point rather than multiple ones. Real-world numerical data from a major automotive electronics company – Bosch Automotive Electronics Portugal – provided context for our modeling explorations. While much of the existing literature discusses how and in which situations safety stock or safety time strategies should be applied, our results suggest that the combined use of both inventory buffers may prove to be a fruitful strategy towards minimizing holding costs while maintaining suitable service levels. Indeed, we observed that the sparsity of the material requirements plan as well as the levels of demand and supply variability prove to be important indicators when dimensioning safety inventory buffers. Particularly, the results derived from applying our model to a comprehensive set of components seem to favor the combination of safety stock and safety time for high levels of sparsity of material requirements plan (i.e., lower delivery frequencies). On the other hand, as the sparsity level decreases (i.e., higher delivery frequencies), we find that the need to use safety stock/time buffers is reduced. Additionally, our results also indicate that variations in demand can influence the decision of which buffer method is the most appropriate. In this case, our findings suggest that a combination of both strategies is recommended whenever demand variation increases. In contrast, in the case of overestimation of demand, the adoption of buffer stock seems to prevail over the use of buffer time, especially for A-type components. With regards to variations in supply, we also found a modest trend to combine safety stock with safety time whenever the magnitude of the supplier delays increases.

From a managerial perspective, our results provided evidence that the joint optimization of safety inventory buffers should not be overlooked when optimizing inventory management.

For added confidence and validation, we developed a flexible decision support system that, by embedding the proposed simulation-based optimization approach, facilitates its implementation and use in real-world supply chain contexts. The architecture of such a system was designed to operate with large-scale data with minimal human-intervention, thereby enhancing its use in real-world environments. By comparing the current safety buffer parameterization adopted by the company with the Pareto-optimal decision suggested by our decision support system, we found that, under certain circumstances, it appears to be more cost-effective to combine safety stock with safety time compared to considering these two inventory buffers independently. In particular, we found that, depending on the ABC class of the components, the joint optimization of safety buffers allowed reductions in holding costs ranging on average from 9.7% to 24.6% and improvements to the company's service level from 1.06% to 13.33%.

However, it is of utmost importance to note that this study is an exploratory research, and our results are based on a particular automotive industry context. As such, more research is needed to demonstrate deeper into the benefits of combining both safety buffers in generalized yet related supply chain contexts. Future research could focus on understanding how slight variations in internal production processes impact the correct dimensioning of safety buffers, as well as to develop modeling approaches to dynamically assign weights to the different criteria involved in choosing a single optimal solution from the Pareto front over different MRP planning horizons. Finally, we emphasize the relevance of correctly modeling supply and demand dynamics without considering common Gaussian approaches just for the sake of statistical simplicity.

Overall, although our modeling approaches have natural limitations and we argue in favor of testing the proposed bi-objective optimization model in other types of industrial contexts, we believe that this paper can be an interesting starting point to trigger further research studies with empirical evidence in this field, which so far has been little explored.

References

Altazin, E., Dauzère-Pérès, S., Ramond, F., and Trefond, S. (2020). A multi-objective optimization-simulation approach for real time rescheduling in dense railway systems. *European Journal of Operational Research*, 286(2):662–672.

Alves, T., Machado, V. H., and Machado, V. C. (2004). Modelling MRP systems under uncertainty: Safety stock versus safety time. In *IIE Annual Conference. Proceedings*, page 1. Institute of Industrial and Systems Engineers (IISE).

Avci, M. G. and Selim, H. (2016). A multi-agent system model for supply chains with lateral preventive transshipments: Application in a multi-national automotive supply chain. *Computers in Industry*, 82:28–39.

Avci, M. G. and Selim, H. (2017). A multi-objective, simulation-based optimization framework for supply chains with premium freights. *Expert Systems with Applications*, 67:95–106.

Avci, M. G. and Selim, H. (2018). A multi-objective simulation-based optimization approach for inventory replenishment problem with premium freights in convergent supply chains. *Omega*, 80:153–165.

Axsäter, S. (2006). *Inventory Control*. Springer, 2nd. edition.

Babai, M. Z., Dai, Y., Li, Q., Syntetos, A., and Wang, X. (2022). Forecasting of lead-time demand variance: Implications for safety stock calculations. *European Journal of Operational Research*, 296(3):846–861.

Bahroun, Z. and Belgacem, N. (2019). Determination of dynamic safety stocks for cyclic production schedules. *Operations Management Research*, 12(1):62–93.

Barbosa-Póvoa, A. P., da Silva, C., and Carvalho, A. (2018). Opportunities and challenges in sustainable supply chain: An operations research perspective. *European Journal of Operational Research*, 268(2):399–431.

Ben-Ammar, O., Bettayeb, B., and Dolgui, A. (2019). Optimization of multi-period supply planning under stochastic lead times and a dynamic demand. *International Journal of Production Economics*, 218:106–117.

Ben-Ammar, O., Dolgui, A., and Wu, D. D. (2018). Planned lead times optimization for multi-level assembly systems under uncertainties. *Omega*, 78:39–56.

Benitez-Hidalgo, A., Nebro, A. J., Garcia-Nieto, J., Oregi, I., and Del Ser, J. (2019). jmetalpy: A python framework for multi-objective optimization with metaheuristics. *Swarm and Evolutionary Computation*, 51:100598.

Beutel, A.-L. and Minner, S. (2012). Safety stock planning under causal demand forecasting. *International Journal of Production Economics*, 140(2):637–645.

Bogataj, D., Aver, B., and Bogataj, M. (2016). Supply chain risk at simultaneous robust perturbations. *International Journal of Production Economics*, 181:68–78.

Boute, R. N., Disney, S. M., Lambrecht, M. R., and Van Houdt, B. (2014). Coordinating lead times and safety stocks under autocorrelated demand. *European Journal of Operational Research*, 232(1):52–63.

Braglia, M., Castellano, D., and Frosolini, M. (2016). A novel approach to safety stock management in a coordinated supply chain with controllable lead time using present value. *Applied Stochastic Models in Business and Industry*, 32(1):99–112.

Buzacott, J. A. and Shanthikumar, J. (1994). Safety stock versus safety time in MRP controlled production systems. *Management Science*, 40(12):1678–1689.

Caridi, M. and Cigolini, R. (2002). Improving materials management effectiveness: A step towards agile enterprise. *International Journal of Physical Distribution & Logistics Management*.

Carlson, R. C. and Yano, C. A. (1986). Safety stocks in mrp?systems with emergency setups for components. *Management Science*, 32(4):403–412.

Chakraborty, T. and Giri, B. (2012). Joint determination of optimal safety stocks and production policy for an imperfect production system. *Applied Mathematical Modelling*, 36(2):712–722.

Chang, C. A. (1985). The interchangeability of safety stocks and safety lead time. *Journal of Operations Management*, 6(1):35–42.

Cheung, K. L. and Hausman, W. H. (1997). Joint determination of preventive maintenance and safety stocks in an unreliable production environment. *Naval Research Logistics*, 44(3):257–272.

Collier, D. A. (1982). Aggregate safety stock levels and component part commonality. *Management Science*, 28(11):1296–1303.

Deb, K. (2000). An efficient constraint handling method for genetic algorithms. *Computer Methods in Applied Mechanics and Engineering*, 186(2-4):311–338.

Deb, K. (2005). *Multi-objective optimization using evolutionary algorithms*. Wiley.

Deb, K., Agrawal, R. B., et al. (1995). Simulated binary crossover for continuous search space. *Complex systems*, 9(2):115–148.

Deb, K., Pratap, A., Agarwal, S., and Meyarivan, T. (2002). A fast and elitist multiobjective genetic algorithm: NSGA-II. *IEEE Transactions on Evolutionary Computation*, 6(2):182–197.

Den Boer, A. A. (1994). Decision support for material planners. *Production Planning & Control*, 5(3):253–257.

Deng, Q., Santos, B. F., and Verhagen, W. J. (2021). A novel decision support system for optimizing aircraft maintenance check schedule and task allocation. *Decision Support Systems*, 146:113545.

Dickersbach, J. T. (2009). Stock and safety stock. In *Supply Chain Management with SAP APO*, pages 415–419. Springer.

Durillo, J. J., Nebro, A. J., Coello, C. A. C., García-Nieto, J., Luna, F., and Alba, E. (2010). A study of multiobjective metaheuristics when solving parameter scalable problems. *IEEE Transactions on Evolutionary Computation*, 14(4):618–635.

Eppen, G. D. and Martin, R. K. (1988). Determining safety stock in the presence of stochastic lead time and demand. *Management Science*, 34(11):1380–1390.

Etienne, E. C. (1987). Choosing optimal buffering strategies for dealing with uncertainty in MRP. *Journal of Operations Management*, 7(1-2):107–120.

Fildes, R. and Kingsman, B. (2011). Incorporating demand uncertainty and forecast error in supply chain planning models. *Journal of the Operational Research Society*, 62(3):483–500.

Fisher, M. L. (1997). What is the right supply chain for your product? *Harvard Business Review*, 75:105–117.

Fonseca, C. M., Paquete, L., and López-Ibánez, M. (2006). An improved dimension-sweep algorithm for the hypervolume indicator. In *IEEE International Conference on Evolutionary Computation*, pages 1157–1163. IEEE.

Glock, C. H. and Kim, T. (2016). Safety measures in the joint economic lot size model with returnable transport items. *International Journal of Production Economics*, 181:24–33.

Gonçalves, J. N., Carvalho, M. S., and Cortez, P. (2020). Operations research models and methods for safety stock determination: A review. *Operations Research Perspectives*, 7:100164.

Grasso, E. T. and Taylor III, B. W. (1984). A simulation-based experimental investigation of supply/timing uncertainty in MRP systems. *International Journal of Production Research*, 22(3):485–497.

Graves, S. C. and Schoenmeyr, T. (2016). Strategic safety-stock placement in supply chains with capacity constraints. *Manufacturing & Service Operations Management*, 18(3):445–460.

Graves, S. C. and Willems, S. P. (2000). Optimizing strategic safety stock placement in supply chains. *Manufacturing & Service Operations Management*, 2(1):68–83.

Graves, S. C. and Willems, S. P. (2003). Supply chain design: safety stock placement and supply chain configuration. *Handbooks in Operations Research and Management Science*, 11:95–132.

Grubbström, R. W. (1998). A net present value approach to safety stocks in planned production. *International Journal of Production Economics*, 56:213–229.

Guide Jr, V. D. R. and Srivastava, R. (2000). A review of techniques for buffering against uncertainty with MRP systems. *Production Planning & Control*, 11(3):223–233.

Hollander, M., Wolfe, D. A., and Chicken, E. (2013). *Nonparametric statistical methods*, volume 751. John Wiley & Sons.

Inderfurth, K. and Minner, S. (1998). Safety stocks in multi-stage inventory systems under different service measures. *European Journal of Operational Research*, 106(1):57–73.

Inderfurth, K. and Vogelgesang, S. (2013). Concepts for safety stock determination under stochastic demand and different types of random production yield. *European Journal of Operational Research*, 224(2):293–301.

Kanet, J. J., Gorman, M. F., and Stößlein, M. (2010). Dynamic planned safety stocks in supply networks. *International Journal of Production Research*, 48(22):6859–6880.

Kania, A., Sipilä, J., Afsar, B., and Miettinen, K. (2021). Interactive multiobjective optimization in lot sizing with safety stock and safety lead time. In *International Conference on Computational Logistics*, pages 208–221. Springer.

Koh, S. C. L., Saad, S. M., and Jones, M. (2002). Uncertainty under MRP-planned manufacture: review and categorization. *International Journal of Production Research*, 40(10):2399–2421.

Lambert, D. M. (2008). *Supply chain management: processes, partnerships, performance*. Supply Chain Management Inst.

Lambert, D. M. and Cooper, M. C. (2000). Issues in supply chain management. *Industrial Marketing Management*, 29(1):65–83.

Lambrecht, M. R., Muckstadt, J. A., and Luyten, R. (1984). Protective stocks in multi-stage production systems. *International Journal of Production Research*, 22(6):1001–1025.

Li, H. and Zhang, Q. (2008). Multiobjective optimization problems with complicated pareto sets, MOEA/D and NSGA-II. *IEEE Transactions on Evolutionary Computation*, 13(2):284–302.

Louly, M.-A. and Dolgui, A. (2013). Optimal MRP parameters for a single item inventory with random replenishment lead time, POQ policy and service level constraint. *International Journal of Production Economics*, 143(1):35–40.

Louly, M.-A., Dolgui, A., and Hnaien, F. (2008). Supply planning for single-level assembly system with stochastic component delivery times and service-level constraint. *International Journal of Production Economics*, 115(1):236–247.

Louly, M.-A. O. and Dolgui, A. (2009). Calculating safety stocks for assembly systems with random component procurement lead times: A branch and bound algorithm. *European Journal of Operational Research*, 199(3):723–731.

Lu, H., Wang, H., Xie, Y., and Li, H. (2016). Construction material safety-stock determination under nonstationary stochastic demand and random supply yield. *IEEE Transactions on Engineering Management*, 63(2):201–212.

Minner, S. (2001). Strategic safety stocks in reverse logistics supply chains. *International Journal of Production Economics*, 71(1-3):417–428.

Minner, S. (2012). *Strategic safety stocks in supply chains*, volume 490. Springer Science & Business Media.

Molinder, A. (1997). Joint optimization of lot-sizes, safety stocks and safety lead times in an MRP system. *International Journal of Production Research*, 35(4):983–994.

Nebro, A. J., Durillo, J. J., Garcia-Nieto, J., Coello, C. C., Luna, F., and Alba, E. (2009a). SMPSO: A new PSO-based metaheuristic for multi-objective optimization. In *2009 IEEE Symposium on Computational Intelligence in Multi-Criteria Decision-Making (MCDM)*, pages 66–73.

Nebro, A. J., Durillo, J. J., Luna, F., Dorronsoro, B., and Alba, E. (2009b). Mocell: A cellular genetic algorithm for multiobjective optimization. *International Journal of Intelligent Systems*, 24(7):726–746.

Nenni, M. E. and Schiraldi, M. M. (2013). Validating virtual safety stock effectiveness through simulation. *International Journal of Engineering Business Management*, 5:41.

Persona, A., Battini, D., Manzini, R., and Pareschi, A. (2007). Optimal safety stock levels of subassemblies and manufacturing components. *International Journal of Production Economics*, 110(1-2):147–159.

Petropoulos, F., Wang, X., and Disney, S. M. (2019). The inventory performance of forecasting methods: Evidence from the M3 competition data. *International Journal of Forecasting*, 35(1):251–265.

Prak, D., Teunter, R., and Syntetos, A. (2017). On the calculation of safety stocks when demand is forecasted. *European Journal of Operational Research*, 256(2):454–461.

Ptak, C. and Smith, C. (2016). *Demand Driven Material Requirements Planning (DDMRP)*. New York: Industrial Press.

Roy, R. K. (2001). *Design of experiments using the Taguchi approach: 16 steps to product and process improvement*. John Wiley & Sons.

Ruiz-Torres, A. J. and Mahmoodi, F. (2010). Safety stock determination based on parametric lead time and demand information. *International Journal of Production Research*, 48(10):2841–2857.

Russell, J. (2014). *Getting started with Impala: interactive SQL for Apache Hadoop*. O'Reilly Media, Inc.

Saoud, P., Kourentzes, N., and Boylan, J. (2021). Approximations for the lead time variance: A forecasting and inventory evaluation. *Available at SSRN 3822270*.

Sarkar, B., Sana, S. S., and Chaudhuri, K. (2010). Optimal reliability, production lotsize and safety stock: An economic manufacturing quantity model. *International Journal of Management Science and Engineering Management*, 5(3):192–202.

Sato, R. and Tsai, T. (2004). Agile production planning and control with advance notification to change schedule. *International Journal of Production Research*, 42(2):321–336.

Shvachko, K., Kuang, H., Radia, S., and Chansler, R. (2010). The hadoop distributed file system. In *2010 IEEE 26th Symposium on Mass Storage Systems and Technologies (MSST)*, pages 1–10. IEEE.

Silver, E. A., Pyke, D. F., and Thomas, D. J. (2016). *Inventory and production management in supply chains*. CRC Press.

Simchi-Levi, D., Kaminsky, P., and Simchi-Levi, E. (2007). *Designing and managing the supply chain: concepts, strategies and case studies*. McGraw-Hill, Singapore, 3rd. edition.

Syntetos, A. A., Babai, Z., Boylan, J. E., Kolassa, S., and Nikolopoulos, K. (2016). Supply chain forecasting: Theory, practice, their gap and the future. *European Journal of Operational Research*, 252(1):1–26.

Thusoo, A., Sarma, J. S., Jain, N., Shao, Z., Chakka, P., Anthony, S., Liu, H., Wyckoff, P., and Murthy, R. (2009). Hive: a warehousing solution over a map-reduce framework. *Proceedings of the VLDB Endowment*, 2(2):1626–1629.

Ting, K. and Cecho, J. J. (2013). *Apache sqoop cookbook: Unlocking hadoop for your relational database*. O'Reilly Media, Inc.

Trapero, J. R., Cardos, M., and Kourentzes, N. (2019a). Empirical safety stock estimation based on kernel and garch models. *Omega*, 84:199–211.

Trapero, J. R., Cardós, M., and Kourentzes, N. (2019b). Quantile forecast optimal combination to enhance safety stock estimation. *International Journal of Forecasting*, 35(1):239–250.

Turan, H. H., Elsawah, S., and Ryan, M. J. (2020). A long-term fleet renewal problem under uncertainty: A simulation-based optimization approach. *Expert Systems with Applications*, 145:113158.

Van Kampen, T. J., Van Donk, D. P., and Van Der Zee, D.-J. (2010). Safety stock or safety lead time: coping with unreliability in demand and supply. *International Journal of Production Research*, 48(24):7463–7481.

Wang, G., Gunasekaran, A., Ngai, E. W., and Papadopoulos, T. (2016). Big data analytics in logistics and supply chain management: Certain investigations for research and applications. *International Journal of Production Economics*, 176:98–110.

Whybark, D. C. and Williams, J. G. (1976). Material requirements planning under uncertainty. *Decision Sciences*, 7(4):595–606.

You, F. and Grossmann, I. E. (2008). Design of responsive supply chains under demand uncertainty. *Computers & Chemical Engineering*, 32(12):3090–3111.

Zaharia, M., Xin, R. S., Wendell, P., Das, T., Armbrust, M., Dave, A., Meng, X., Rosen, J., Venkataraman, S., Franklin, M. J., et al. (2016). Apache spark: a unified engine for big data processing. *Communications of the ACM*, 59(11):56–65.

Zhong, R. Y., Newman, S. T., Huang, G. Q., and Lan, S. (2016). Big data for supply chain management in the service and manufacturing sectors: Challenges, opportunities, and future perspectives. *Computers & Industrial Engineering*, 101:572–591.

Zitzler, E., Laumanns, M., and Thiele, L. (2001). SPEA2: improving the strength pareto evolutionary algorithm. *TIK-report*, 103.

Zitzler, E. and Thiele, L. (1998). Multiobjective optimization using evolutionary algorithms—a comparative case study. In *International conference on parallel problem solving from nature*, pages 292–301. Springer.

Zitzler, E. and Thiele, L. (1999). Multiobjective evolutionary algorithms: a comparative case study and the strength pareto approach. *IEEE Transactions on Evolutionary Computation*, 3(4):257–271.