



HAL
open science

Material Requirements Planning Under Demand Uncertainty Using Stochastic Optimization

Simon Thevenin, Yossiri Adulyasak, Jean-françois Cordeau

► **To cite this version:**

Simon Thevenin, Yossiri Adulyasak, Jean-françois Cordeau. Material Requirements Planning Under Demand Uncertainty Using Stochastic Optimization. Production and Operations Management, 2020, 10.1111/poms.13277 . hal-03052916

HAL Id: hal-03052916

<https://hal.science/hal-03052916v1>

Submitted on 14 Feb 2022

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution - NonCommercial 4.0 International License

Material Requirements Planning Under Demand Uncertainty Using Stochastic Optimization

Simon Thevenin

IMT Atlantique, LS2N-CNRS, 44307 Nantes, France;

Corresponding author; `simon.thevenin@imt-atlantique.fr`; phone:+332 51 85 82 26

Yossiri Adulyasak and Jean-François Cordeau

HEC Montréal and GERAD, Montréal, QC H3T 2A7, Canada;

`yossiri.adulyasak@hec.ca`; phone:+1514 340-7029; fax:+1514 340-6834

`jean-francois.cordeau@hec.ca`; phone:+1514 340-6278; fax:+1514 340-6834

Abstract

Material Requirements Planning (MRP), a core component of enterprise resource planning (ERP) systems, is widely used by manufacturers to determine the production lot sizes of components. These lot sizes are typically computed based on deterministic and dynamic demand assumptions, while safety stocks, which hedge against demand uncertainty, are determined independently based on different assumptions. As the lot sizes and safety stocks are not determined simultaneously, sub-optimal decisions are often used in practice. The critical impact of inventories and service levels in manufacturing motivates the study of stochastic optimization methods for MRP. In this paper, we investigate stochastic optimization methods for MRP systems under demand uncertainty. A two-stage and a multi-stage model are proposed to deal with the static-static and static-dynamic decision frameworks, respectively. We first derive structural properties of the two-stage and multi-stage models to provide insights on the differences between the plans created with these two models. As multi-stage stochastic programs are not convenient in real-world applications, several practical enhancements are proposed. First, to address scalability issues, we employ heuristics in combination with advanced sampling methods. Second, to allow real-time static-dynamic decisions, we derive a policy from the solution of the multi-stage model.

Third, to deal with the dynamic-dynamic decision framework, we employ a rolling horizon implementation. The effectiveness and performance of stochastic optimization for MRP are validated by numerical experiments, which demonstrate that the stochastic optimization approaches have the potential to generate significant cost savings compared to traditional methods for production planning and safety stocks determination.

Keywords: Material requirements planning; stochastic optimization; lot-sizing; uncertain demand

History: Received: July 2019; accepted: September 2020 by Chelliah Sriskandarajah after three revisions.

1 Introduction

The global market for enterprise resource planning (ERP) systems was worth \$39.65 billion in 2019 and is expected to reach \$65.20 billion by 2025 (Mordor Intelligence 2020). At the core of these ERP systems is the Material Requirements Planning (MRP) logic. Given the demands for end items, MRP systems compute the sizes of the lots to produce (or to order) for each component in each period. These calculations are based on the bills of materials (BOMs), which indicate the hierarchy of components (i.e., the structure of the product, and the number of components required to produce each end item or subsequent component). MRP systems are often used in an uncertain environment. However, in this context, classical MRP systems create inadequate plans because their computations utilize demand forecasts considered as deterministic (Chevreux et al. 2018). Although these systems can redetermine the production plan relatively quickly based on recently updated demand information, such reactive changes are far from optimal. Indeed, the new production plan is subject to the previously implemented plan, and it lacks a proactive vision to deal with the uncertainty in future demand. To alleviate this issue, practitioners often attempt to improve the demand forecasting accuracy. This countermeasure alone is far from being perfect as one cannot practically fully eliminate the forecasting errors. Hence, in addition to the effort to improve demand forecasting accuracy, it is inevitable for manufacturing companies to rely on an approach that can properly account for stochastic demand (Cecere 2015). Consequently, and as a step in this direction, this paper investigates the use of stochastic optimization for MRP systems.

The problem solved by MRP systems is a multi-echelon, multi-item, lot-sizing problem (Tem-

pelmeier and Derstroof 1996). Early implementations of MRP solved this problem item by item with heuristic lot-sizing rules based on simple logic. For instance, the lot-for-lot rule sets the production quantities to the requirements of each period. As these rules perform poorly in the presence of shared resource capacities, recent implementations of MRP solve the lot-sizing problem with a mixed-integer linear program (MILP). The MILP approach is flexible enough to model multiple extensions and constraints. These problems are commonly solved under the assumption of deterministic demand since safety stocks are calculated separately to hedge against uncertainty. However, the existing safety stock computation methods for multi-echelon production systems do not consider the key decision components of MRP systems. In fact, the applicable methods (e.g., Graves and Willems 2008) are often designed for base stock policies. As the application of a base stock policy follows simple rules, the level of stock can be written as a function of the stochastic demand, and the safety stocks are computed to meet a given service level. This approach is not suitable for advanced MRP systems, where the lot sizes must be determined in a complex environment with multiple echelons, setup costs, capacities, etc. (see Section 2). Consequently, sub-optimal safety stock levels (not determined in conjunction with the lot sizes) are often used in practice. These safety stock levels are computed either manually, at the master production schedule (MPS) level, or under the assumption of a base stock policy (e.g., Graves and Willems 2008).

We consider two decision frameworks, referred to as *static-static*, and *static-dynamic* in the lot-sizing literature (e.g., Tempelmeier 2013), and we model the production planning problem under these decision frameworks using stochastic programming formulations. The static-static environment is encountered when a frozen period is considered. That is, the production quantities and setups are decided at time zero (the current period prior to demand realization) and fixed for the entire time horizon. The static-dynamic situation occurs when setups are linked to long-term decisions and must be fixed, whereas production quantities can be adjusted in each period. More precisely, the setups are decided at time zero for the entire planning horizon, whereas the production quantities for period $t+1$ are decided after having observed the demands of period t . Examples of long-term decisions linked with the setups include the planning of secondary resources, such as the models in 3D printing, the molds in injection moldings, or the technicians who set up the production lines. Indeed, planning these secondary resources requires knowledge of the items produced in each period.

The contributions of this paper are threefold. First, to tackle demand uncertainty in MRP, we design scenario based stochastic optimization approaches for the *multi-echelon multi-item capacitated lot-sizing problem* (MMCLP) with lead times and stochastic demand. As these methods account for the demand’s stochasticity in the lot-sizing model, they remove the boundary between the lot sizes and safety stocks computations. Compared to the commonly used decomposition of the safety stocks computation and lot sizing, these techniques provide significant cost savings and are highly flexible. Second, the scalability problems arising with conventional stochastic programming approaches are addressed here with efficient heuristics, advanced scenario sampling methods, two real-time execution policies derived from the solution of the stochastic models (denoted S-Policy and Q-policy), as well as a rolling horizon solution framework to deal with long planning horizons and to deal with the dynamic-dynamic decision framework (Bookbinder and Tan 1988). This dynamic-dynamic variant is the most challenging one. To the best of our knowledge, no efficient method has been proposed for this variant of the stochastic multi-stage lot-sizing model. Third, we present managerial insights on the stochastic optimization approaches derived from the structural results, and we validate their effectiveness and performance through numerical experiments. Compared to common methods in practice and the methods in the literature for production planning and safety stocks determination, the stochastic optimization approaches based on two-stage and multi-stage stochastic programming models show a strong potential to generate significant cost savings. Despite additional complexity in the multi-stage models, their adaptability allows the manufacturer to gain additional benefits in a dynamic decision environment. To the best of our knowledge, this paper is also the first to present an extensive numerical simulation of stochastic optimization approaches for practical multi-echelon MRP settings. The stochastic approaches proposed in this paper scale well, and MRP software developers can directly implement these efficient methods. Since the MRP logic is also used in Distribution Resources Planning (DRP) systems, the solution approaches presented in this study can also be applied directly in DRP systems under demand uncertainty.

The paper presents extensive numerical experiments to validate the value of stochastic optimization in MRP software. The stochastic optimization approaches are compared with classical methods, namely, the deterministic mathematical model, and lot-sizing rules (lot-for-lot, economic order quantity, economic order period, Silver-Meal) equipped with safety stocks. The methods are simulated with a large number of scenarios using well-known academic benchmarks. As the

proposed methods require no assumptions on the demand’s probability distribution, three types of distributions are considered in the experiments: (1) to model the classical use of MRP, the demands follow *non-stationary* Normal distributions; (2) to model *slow moving* items, the demands follow Poisson distributions; (3) to model items with lumpy demands, the demands follow zero-inflated Poisson distributions. Besides the static-static and static-dynamic decision frameworks, the considered models are also evaluated with a rolling horizon simulation in the *dynamic-dynamic* environment, where the plan is re-optimized in each period. Table 1 gives the model, solution approaches, and evaluation methodology for each decision framework. These experiments complement the theoretical insights gained from the structural properties of the problem. It is evident that traditional methods that determine lot sizes and safety stocks separately yield sub-optimal plans which result in higher overall costs. Therefore, manufacturing companies should consider adopting stochastic optimization in MRP systems to determine more cost-effective production plans which can be executed in a dynamic fashion.

Decision framework	Model	Solution approaches	Evaluation
Static-static	Two-stage	Two-stage	Observe (No-replanning)
Static-dynamic	Multi-stage	Two-stage heuristic, Multi-stage, Fix-and-optimize, S-Policy	Re-solve
Dynamic-dynamic	-	Two-stage heuristic, Multi-stage heuristic, Fix-and-optimize, S-Policy, Q-Policy	Rolling horizon simulation

Table 1: Considered decision framework.

This paper is organized as follows. Section 2 gives a review of previous work on stochastic MRP and stochastic multi-echelon lot-sizing problems. Section 3 presents stochastic optimization models for MRP under demand uncertainty along with the structural analyses. Section 4 presents the scenario sampling approaches, the proposed heuristics, and the order-up-to-level policy. Section 5 presents the methods used to benchmark stochastic optimization approaches and the simulation framework. Section 6 reports the experimental results. Finally, the conclusion follows in Section 8.

2 Literature Review

Mathematical models for MRP have mostly been studied in a deterministic context (e.g., Zahorik et al. 1984, Billington et al. 1983, Clark and Armentano 1995). However, in practice, MRP systems are subject to diverse forms of uncertainty: demand, lead times, production yield, production

capacity, among others (Guide and Srivastava 2000, Dolgui and Prodhon 2007). This literature review focuses on demand uncertainty in multi-echelon production systems, where the following three topics are successively covered: simulation studies on MRP in a stochastic demand context, safety stocks for MRP, and stochastic optimization approaches for multi-echelon MRP.

Several authors (e.g., Bai et al. 2002, Zhao and Lee 1993, Zhao et al. 2001, Enns 2002, Kadipasaoglu and Sridharan 1995, Ho and Ireland 1998) evaluate by simulation the impact of demand uncertainty on MRP systems for multi-echelon production problems. These studies also evaluate how the parameters (safety stocks, safety lead times, re-planning frequencies, frozen periods, sizing rules) of classical MRP systems protect against demand uncertainty. The parameters considered by each of these papers are summarized in Table 2. The main conclusions of these works are the following. First, demand uncertainty and forecast errors have a significant impact on the costs and service levels. Second, safety stocks and safety lead times (i.e., considering buffer lead times in addition to the expected lead times) are efficient ways to protect against stochastic demand, but the choice of one versus the other depends on the considered system. In addition, most studies (e.g., Lagodimos and Anderson 1993, Bai et al. 2002, Zhao et al. 2001, Boulaksil 2016) advise to place safety stocks at the end item level, but some studies disagree. For instance, Carlson and Yano (1986) suggest holding some safety stocks for components with large setup costs. Third, frequent re-planning with classical MRP systems is undesirable, because MRP systems are prone to nervousness (i.e., a minor change in the data leads to large modifications of the plan), and users tend not to trust a nervous system (Blackburn et al. 1985). Kadipasaoglu and Sridharan (1995) and Zhao and Lee (1993) showed that frequent re-planning leads to larger costs than infrequent re-planning and that freezing a part of the master production schedule is the most efficient way to reduce nervousness.

Paper	Counter measures
Bai et al. (2002)	Frozen period, lot-sizing, safety stocks, planning horizon
Zhao and Lee (1993)	Frozen period, planning horizon, re-planning frequency
Zhao et al. (2001)	Safety stocks
Enns (2002)	Safety stocks, safety lead times, lot-sizing
Kadipasaoglu and Sridharan (1995)	Frozen period, safety stocks, lot-sizing
Ho and Ireland (1998)	Lot-sizing

Table 2: Previous studies on multi-echelon MRP systems with stochastic demand.

Although multiple studies suggest using safety stocks in MRP systems with demand uncertainty,

to the best of our knowledge, no analytical method exists to directly determine the safety stocks in an MRP environment. In fact, the few works on safety stocks for MRP systems focus on special cases because the behavior of an MRP system is hard to model analytically (Benton 1991). For instance, Lagodimos and Anderson (1993) propose a safety stocks computation approach for an MRP system with a lot-for-lot policy, constant demand, serial network, and no holding cost. Zijm and Van Houtum (1994) consider an MRP system with an order-up-to-level policy in assembly systems. Inderfurth (2009) studies a single-echelon MRP system with critical stock policy, where the production quantities are computed to bring the inventory levels above some critical thresholds. Other works on safety stocks for multi-echelon systems with non-stationary and uncertain demand (e.g., Inderfurth and Minner 1998, Graves and Willems 2008, Graves and Schoenmeyr 2016) focus on base stock policies. These approaches do not directly apply to the context of MRP, where the lot-sizing decisions have a significant impact on the risk of shortage.

Instead of computing the safety stock levels analytically, Benton (1991) and Boulaksil (2016) propose to use simulation methods. The lot-sizing problem is first solved based on the expected demand. Then, a simulation is run, and the production quantities are adjusted to meet the desired service level. In the same vein, Sali and Giard (2015) revised the lot-for-lot rule to deal with uncertain demand. In their approach, lot sizes are computed to achieve the desired service level according to the cumulative distribution of the projected inventory levels.

Stochastic optimization models remove the need for safety stocks because they account implicitly for the demand's probability distributions. In other words, the computation of safety stock levels and lot sizes are no longer isolated. Tempelmeier (2013) and Aloulou et al. (2014) review stochastic optimization methods for lot-sizing problems with uncertain demand. Most studies (e.g., Brandimarte 2006) consider a single-echelon production system. To the best of our knowledge, the only work considering stochastic optimization for multi-echelon systems is presented in Grubbström and Wang (2003). The authors propose a dynamic programming approach to minimize the net present value in the capacitated multi-echelon lot-sizing problem with stochastic demand but without lead times. However, stochastic optimization approaches have been proposed for related problems such as supply chain management problems (e.g., Lin and Uzsoy 2016). In addition, other approaches than stochastic optimization have been proposed for the single item lot-sizing problem with stochastic demand such as distributionally robust optimization (e.g., Zhang et al. 2016).

Our work differs from the above literature in several aspects. First, to the best of our knowledge, this paper is the first to investigate scenario based multi-stage stochastic optimization for capacitated multi-echelon MRP systems with lead times and stochastic dynamic demand. These methods are useful for practitioners since they remove the boundary between safety stocks and lot sizes computations in MRP systems. Consequently, the proposed approaches remove the problems associated with safety stocks computations in MRP systems. In addition, we compare the performance of a two-stage and a multi-stage formulation in static-static, static-dynamic, and dynamic-dynamic decision framework. To alleviate the scalability issues of stochastic optimization approaches, advanced sampling methods are considered, as well as a fix-and-optimize heuristic, and an execution policy. Finally, computationally intensive simulations are performed to compare (in terms of costs and key performance indicators — KPIs) stochastic models with classical approaches such as lot-sizing rules, and deterministic models. The results show that stochastic optimization leads to significant cost savings in MRP systems. In addition, the proposed methods are efficient and scalable.

3 Stochastic MRP Models under Demand Uncertainty

This section describes the considered problem (Section 3.1) and the proposed stochastic optimization formulations modeling the static-static (Section 3.2) and static-dynamic (Section 3.3) decision frameworks. To illustrate the MMCLP, the Electronic Companion provides a small example.

3.1 MRP Under Demand Uncertainty

MRP systems are used to determine the production quantity Q_{it} for each item i in a set \mathcal{I} and for each period t in the time horizon $\mathcal{H} = \{1, \dots, T\}$. The inputs of the model include BOM, lead times, probability distributions of the demands, and production capacities. We denote by $\mathcal{I} = \mathcal{I}_e \cup \mathcal{I}_c$ the set of items, where \mathcal{I}_e and \mathcal{I}_c are the sets of end items and of components, respectively. We assume (without loss of generality) that components have no external demand, whereas the probability distribution \tilde{D}_{it} of the demand is known for each end item $i \in \mathcal{I}_e$ and each period t . The BOM gives the hierarchy of components required for each end item, that is, the number of units R_{ij} of item i required to produce one unit of j . In addition, each item i has a lead time L_i , i.e., the production quantity Q_{it} is available for the next production step in period $t + L_i$. Finally, the

production plan must respect the capacity C_k of each resource, given the resource consumption K_{ik} per unit of item i for each resource k in the set \mathcal{K} .

For each item i , inventory holding costs (h_i), fixed setup costs (s_i), and unit production costs (v_i) are considered. In addition, backlog costs (b_i) and lost sale costs (e_i) are considered for the end items. The unmet demand of item i in period t is backlogged (it can be fulfilled in subsequent periods), but a penalty b_i is incurred in each period for each unit of backlog.

Note that in the deterministic lot-sizing problem, we typically impose a constraint to ensure that the full demand is met within the horizon. In the case of stochastic demand, however, this constraint leads to overly conservative decisions since it requires satisfying the worst-case demand among all scenarios. The inclusion of a lost sale penalty e_i at the end of the horizon, as commonly used in the literature (e.g., Absi et al. 2011), allows the model to make the optimal trade-off while properly accounting for demand uncertainty. Finally, there is no constraint on the ordering quantities of raw materials (we assume that suppliers have an infinite capacity).

3.2 Two-Stage Stochastic MRP Model for the Static-Static Framework

The stochastic formulations of MMCLP are based on the set Ω of all possible demand scenarios. Given the probability p_ω of each scenario ω (with $p_\omega > 0$ and $\sum_{\omega \in \Omega} p_\omega = 1$), the problem is to find the solution with the minimum expected total cost. The static-static decision framework can be represented by a two-stage stochastic optimization model. The first-stage variables correspond to the decisions made in period 0 (while the demand is unknown). In the static-static MMCLP, these decisions are the setup Y_{it} (variable equal to 1 if there is a setup, and 0 otherwise) and quantity Q_{it} for item i in period t . Second stage variables correspond to the inventory I_{it}^ω and backlog level B_{it}^ω of item i at the end of period t , observed after the realization of the demand to compute the cost for each scenario ω . B_{iT}^ω indicates the total remaining backlog quantity of item i for scenario ω at the end of the horizon, and this quantity can be interpreted as the lost sale. As the setup and quantity decisions are made before observing the demand, they are the same for all scenarios. On the contrary, the second stage variables can be different for each scenario. The deterministic capacitated multi-echelon lot-sizing problem is NP-hard, as it extends the capacitated lot-sizing problem which is itself NP-hard (Bitran and Yanasse 1982). As the stochastic versions extend the problem with multiple scenarios, it is also NP-hard.

The problem can be formulated as the following MILP:

$$\min \sum_{\omega \in \Omega} p_{\omega} \left(\sum_{t \in \mathcal{H}} \sum_{i \in \mathcal{I}} (h_i I_{it}^{\omega} + s_i Y_{it} + v_i Q_{it}) + \sum_{i \in \mathcal{I}_e} \left(\sum_{t=1}^{t=T-1} b_i B_{it}^{\omega} + e_i B_{iT}^{\omega} \right) \right) \quad (1)$$

$$\text{s.t.} \quad \sum_{\tau=1}^{t-L_i} Q_{i\tau} + I_{i0} - \sum_{\tau=1}^t D_{i\tau}^{\omega} - I_{it}^{\omega} + B_{it}^{\omega} = 0 \quad i \in \mathcal{I}_e, \quad t \in \mathcal{H}, \quad \omega \in \Omega \quad (2)$$

$$\sum_{\tau=1}^{t-L_i} Q_{i\tau} + I_{i0} - \sum_{\tau=1}^t \left(\sum_{j \in \mathcal{I}} R_{ij} \cdot Q_{j\tau} \right) - I_{it}^{\omega} = 0 \quad i \in \mathcal{I}_c, \quad t \in \mathcal{H}, \quad \omega \in \Omega \quad (3)$$

$$Q_{it} \leq M_i Y_{it} \quad i \in \mathcal{I}, \quad t \in \mathcal{H} \quad (4)$$

$$\sum_{i \in \mathcal{I}} K_{ik} Q_{it} \leq C_k \quad t \in \mathcal{H}, \quad k \in \mathcal{K} \quad (5)$$

$$B_{it}^{\omega} \geq 0 \quad i \in \mathcal{I}_e, \quad t \in \mathcal{H}, \quad \omega \in \Omega \quad (6)$$

$$I_{it}^{\omega} \geq 0 \quad i \in \mathcal{I}, \quad t \in \mathcal{H}, \quad \omega \in \Omega \quad (7)$$

$$Q_{it} \geq 0 \quad \text{and} \quad Y_{it} \in \{0, 1\} \quad i \in \mathcal{I}, \quad t \in \mathcal{H}. \quad (8)$$

The objective function (9) is the expected total cost over all the scenarios, including inventory costs, setup costs, unit production costs, backlog costs, and end-of-horizon lost sales costs. Constraints (2) set the value of the backlog and inventory quantities for the end items. These values depend on the produced quantities and external demands. Constraints (3) set the inventory levels of components, which depend on the internal demands. Note that backlogs are not allowed for components since they are required for the planned production. Constraints (4) set the variable Y_{it} to 1 if the quantity of item i produced in period t is greater than 0. The value of M_i in constraints (4) is an upper bound of the production quantity of item i . This upper bound can be set to the minimum between the upper bound M_i^1 (defined in Equation (9)) inferred from the demands of item i , and the upper bound M_i^2 (defined in Equation (10)) inferred from the production capacities:

$$M_i^1 = \begin{cases} \max_{\omega \in \Omega} \sum_{t \in \mathcal{H}} D_{it}^{\omega} & \text{if } i \in \mathcal{I}_e \\ \sum_{j \in \mathcal{I}} R_{ij} \cdot M_j^1 & \text{if } i \in \mathcal{I}_c \end{cases} \quad (9) \quad M_i^2 = \min_{k \in \mathcal{K} | K_{ik} > 0} \frac{C_k}{K_{ik}}. \quad (10)$$

The value of M_i^1 states that the production quantity cannot be larger than the maximum total demand. Constraints (5) ensure that production capacities are respected.

3.3 Multi-Stage Stochastic MRP Model for the Static-Dynamic Framework

The static-dynamic environment corresponds to a multi-stage stochastic optimization model. This model is similar to model (1)-(8), but the production quantities are scenario-dependent, and non-anticipativity constraints are included in the model. In the static-dynamic decision framework, the production quantities in period t depend on the realizations of the demands in periods $1, \dots, t-1$. Therefore, the quantity Q_{it}^ω of item i produced in period t depends on the scenario ω . The *non-anticipativity* constraints (11) ensure that identical decisions are made at stage t in all scenarios indistinguishable up to stage t . Accordingly, constraints (11) enforce equal production quantities in period $t+1$ for all scenarios ω with identical demands $D_\omega^{1\dots t}$ in periods 1 to t . Contrarily to the production quantities, the inventory and backlog levels are observed once the demands are known. Therefore, the index t (and not $t+1$) is used in the non-anticipativity constraints:

$$Q_{it+1}^\omega = Q_{it+1}^{\omega'}, \quad I_{it}^\omega = I_{it}^{\omega'}, \quad B_{it}^\omega = B_{it}^{\omega'} \quad \forall i \in \mathcal{I}, t \in \mathcal{H}, \omega, \omega' | D_\omega^{1\dots t} = D_{\omega'}^{1\dots t}. \quad (11)$$

3.4 Structural Properties and Managerial Insights

While the two-stage and multi-stage models correspond to distinct decision frameworks, the differences between the production plans resulting from these models are not obvious. This section highlights the benefits of the multi-stage model (despite its complexity) to encourage its adoption in practice. The main findings can be summarized by stating that the multi-stage model gives more reactivity to the production planners, with the creation of safety stocks of components and the adjustments of the production quantities to the dynamic order decisions. These findings are supported by Propositions 1-3, and their proofs are provided in the Electronic Companion.

Proposition 1 shows that the two-stage model leads to producing the exact amount of components required for the production of end items. On the contrary, the multi-stage model leads to left-overs at the end of the horizon (see Proposition 2), because the multi-stage model takes into account the possibility of keeping components (instead of transforming them to end items) when the demand is low. That is, the stock of components is computed to react appropriately to the observed demand. This behavior of the multi-stage model can be assimilated with the creation of safety stocks at the components level.

Proposition 1. *The optimal solution of the two-stage formulation for the static-static MMCLP has zero inventory of components in the last period, except for the components whose initial stock level is more than the amount required for the production of the end items during the complete planning horizon.*

Proposition 2. *The optimal solution of the multi-stage formulation for the static-dynamic MMCLP can have positive component inventory levels at the end of the planning horizon.*

MRP models plan the delivery of the right amount X_{it} at the right moment t for each product i . In the static-dynamic context, a reduction of the lead times postpones the production quantities decisions. Therefore, more information is available when the order is passed, and the production quantities are adjusted to the revealed demand. Such behavior occurs in the static-dynamic decision framework but not in the static-static. However, a reduction of the lead times in the two-stage models can impact the production quantity due to the capacity and flow conservation constraints. To isolate the effect of the additional information on the sizing of the lots, we consider the special case of a single item with infinite capacity. Proposition 3 shows that a reduction of the lead time does not change the solution in static-static, but does in static-dynamic.

Proposition 3. *For the special case of MMCLP with a single item and infinite capacity, a reduction of the lead time has no impact on the total cost in the static-static framework (the production quantities are only shifted), as long as the initial inventory covers the demand from period 1 to the lead time L . On the contrary, a reduction of the lead time impacts the production quantities and the total cost in the static-dynamic decision framework.*

4 Solution Approaches

This section presents solution approaches to solve the stochastic optimization models for MRP under demand uncertainty in an efficient way. As the complete set Ω of scenarios is usually too large, scenario sampling methods are provided in Section 4.1. Section 4.2 discusses a formulation of the multi-stage model where the non-anticipativity constraints are modeled implicitly, as well as a fix-and-optimize heuristic developed for this formulation. The resulting approaches significantly reduce memory consumption and computation time. In addition, Section 4.2 introduces an order-up-to-level policy conditional on production setups, which is derived from the multi-stage

model. This approach alleviates the cumbersome static-dynamic decision process since it requires no computation in the subsequent decision stages. Finally, Section 4.4 presents a rolling-horizon heuristic to solve the problem in a dynamic-dynamic decision framework.

4.1 Demand Scenario Sampling

Solving the two-stage (resp. multi-stage) model with the set Ω of all possible scenarios leads to the true optimal solution in the static-static (resp. static-dynamic) decision framework. However, Ω is usually large (sometimes infinite), and solving the resulting MILP is often impossible in practice. Consequently, the problem is approximated with samples of scenarios. This section describes the tree structure required to generate the scenarios in the multi-stage model, before exposing three scenario sampling techniques (used for the multi-stage and two-stage model), namely, crude Monte Carlo (CMC), quasi-Monte Carlo (QMC), and randomized quasi-Monte Carlo (RQMC). The Electronic Companion gives more details on the generation of QMC samples as well as the algorithm to sample with CMC, RQMC, and to generate scenario trees.

The representations of the scenarios are different in the two-stage and multi-stage models. In the two-stage model, a scenario is a vector whose components are the demands for each end item in each period. In the multi-stage model, the scenarios are generated with a scenario tree, as shown in Figure 1. Each level of the tree corresponds to a period, the children of a node at level t are possible realizations of the demands in period $t + 1$, and each path in the tree corresponds to a scenario. The scenario tree ensures that the decided production quantities in period t (given the demands in periods $1, \dots, t - 1$) account for the stochastic demands in periods t, \dots, T . Indeed, in a scenario tree, multiple demand realizations for periods t, \dots, T are available for each demand's realization of periods $1, \dots, t - 1$. The scenario tree structure is denoted by $[N_1, \dots, N_T]$, where N_t is the number of branches of the nodes at level t . For instance, the structure of the tree in Figure 1 is $[2, 2]$. The demand realizations are sampled independently at each node of the tree. Consequently, the multi-stage model requires sampling vectors whose components are the demands for a single period (i.e., the dimension of the vectors is the number of end items).

Note that advanced sampling techniques are crucial for the multi-stage model, since the size of the tree is exponential in the number of periods, and only a few demand realizations can be sampled at each node. For more information on scenario generation in multi-stage stochastic optimization,

the interested reader is referred to Dupačová et al. (2003), and Kaut and Wallace (2007).




Figure 1: Example of scenario tree [2, 2]

CMC samples n vectors randomly (according to the probability distribution of the demand) with equal probability $1/n$. The CMC approximation is on average equal to the true optimal expected cost when the number of scenarios is sufficiently large. However, the variance of the approximation is σ^2/n , where σ^2 is the variance of the expected optimal cost. On the other hand, RQMC finds samples leading to approximations with theoretically lower variances than CMC. QMC and RQMC first select a set V_n of vectors in $[0, 1]^d$, where d is the dimension of the sampled vectors, to cover the unit cube evenly. From V_n , the demand vectors are generated using the inverse of the cumulative probability distribution of the demands. Like CMC, all vectors have equal probability $1/n$. Two methods exist to generate V_n , namely, lattice rules and digital nets. Rank-1 lattice rules are used in this paper (higher rank lattices are uncommon in practice). These rules are formally defined as $V_n = \{i \cdot \alpha/n + \delta \bmod 1 \quad \forall i \in 1 \dots n\}$, where α is a generator vector, and δ is a random point to shift the lattice. Using $\delta = 0$ leads to QMC which is a deterministic sampling technique, whereas RQMC uses $\delta > 0$ to generate a random sample. The quality of the lattice is determined by vector α . We generate α with the software Lattice Builder (L’Ecuyer and Munger 2016), which implements multiple algorithms to build good rank-1 lattice rules.

Finally, as demands are integer, a QMC or RQMC sample can contain multiple occurrences of a vector (even if V_n is fully projection-regular). Identical vectors are aggregated into a single one by adding their probabilities. To better control the number of scenarios, the sample size is increased until a predefined number n of different vectors are obtained. Such aggregations are especially useful for lumpy demands, where the probability of having no demand is large.

4.2 Solution Approaches for the Static-Dynamic Environment

As the number of scenarios of the multi-stage model grows exponentially with the number of periods, solving the multi-stage model is very challenging. To reduce the size of the model, the

formulation with *implicit non-anticipativity* uses a single variable for each group of equal variables. More precisely, the variable $Q_{it+1}^{D(1\dots t)}$ replaces the set $\{Q_{it+1}^\omega \mid D_\omega^{1\dots t} = D(1\dots t)\}$ of variables, which represents the quantity produced in period t in different scenarios ω with identical demands $D(1\dots t)$ from periods 1 to t . Similarly, variables $I_{it}^{D(1\dots t)}$ and $B_{it}^{D(1\dots t)}$ replace the sets $\{I_{it}^\omega \mid D_\omega^{1\dots t} = D(1\dots t)\}$ and $\{B_{it}^{\omega'} \mid D_\omega^{1\dots t} = D(1\dots t)\}$, respectively. Also, constraints (1)-(8) are generated for each possible realization of the demands in period 1 to t (and not for each scenario), item, and period. In addition, to further speed up the solution of the problem, the solution of the two-stage model is used as a warm start for the multi-stage model.

Finally, we propose two heuristics for the static-dynamic framework, denoted, the *two-stage* heuristic and the *fix-and-optimize* heuristic. The *two-stage* heuristic simply solves model (1)-(8), and it ignores the dynamic production quantity decisions. Model (1)-(8) corresponds to the multi-stage model with the following additional constraints,

$$Q_{i,t}^\omega = Q_{i,t}^{\omega'} \quad \forall i \in \mathcal{I}, \quad t \in \mathcal{T}, \quad \omega, \omega' \in \Omega.$$

As the two-stage model is more restricted, its solution is a feasible (but not necessarily optimal) solution in the static-dynamic context. The *fix-and-optimize* approach has two steps. The first step determines the setups by solving the two-stage model. The second step solves the multi-stage formulation, but with the setups fixed to the values found in the first step. As the non-fixed variables in the second step (production quantities, inventory levels, backlogs, lost sales) are continuous, the resulting model is a linear program. The Electronic Companion provides the algorithm with the detailed steps to implement the *fix-and-optimize* heuristic.

The *fix-and-optimize* heuristic is more complex than the two-stage heuristic, and thus it requires more computational time. However, if the two-stage and the *fix-and-optimize* heuristics use the perfect set of scenarios, *fix-and-optimize* yields the optimal production quantities for the fixed setups. Providing additional analytical results on the MMCLP is difficult because MMCLP is very generic, but we analyze below the behavior of the two-stage heuristic on the special case of MMCLP with a single item and infinite capacity. Despite being restricted to a special case, these analytical results show that *fix-and-optimize* outperforms the two-stage heuristic. This special case corresponds to the traditional context of MRP at the MPS level (i.e., the level with end-items

only), which are used in the (rough-cut) capacity planning process.

Proposition 4 shows that the two-stage heuristic leads to overproduction when MRP is used at the MPS level. This situation corresponds to the special case of MMCLP with a single item (and thus a single level) and infinite capacity. In the fix-and-optimize heuristic, once step 1 is performed, the setups are fixed to the values of the solution of the two-stage model, and we denote by $\sigma(k)$ the period in which the k^{th} production lot is completed.

Proposition 4. *For the special case of static-dynamic MMCLP with a single item, infinite capacity, and when $\sigma(1)$ equals to the lead time L , the two-stage heuristic leads to a larger or equal production quantity completed in period $\sigma(1)$ compared to fix-and-optimize.*

Proof. See the Electronic Companion. □

When the two-stage and the fix-and-optimize heuristics are used in a rolling horizon framework, only the decisions of stage 0 (i.e., the lots completed in period $\sigma(1) = L$) are implemented. As a consequence, using the two-stage heuristic in a static-dynamic framework can lead to overproduction.

4.3 MRP Execution Policy

The use of the multi-stage model in the static-dynamic decision framework is cumbersome because it requires to re-solve the model in each period. To ease the process, an order-up-to-level policy (denoted S-policy), which is conditional on the production setups, is derived from the solution of the multi-stage model. The S-policy is based on the notion of echelon stock E_{it} , which denotes the total quantity of item i in the system in period t . The echelon stock includes the stock of item i , the components of i in the stocks of downstream items, and the quantities ordered in previous periods but not yet produced. Before applying the policy, the setup Y_{it} of each item i and period t is taken from the solution of the multi-stage model, and the values of the replenishment level S_{it} are inferred as follows. First, for each optimization scenario ω , the replenishment level S_{it}^ω is calculated as $S_{it}^\omega = E_{it}^\omega + Q_{it}^\omega$. To compute S_{it} , we consider only the scenarios where Q_{it}^ω is not constrained by the capacities or flow conservation constraints (but if the quantity is constrained in all the scenarios, then all the scenarios are considered). S_{it} corresponds to the average value of S_{it}^ω in these scenarios. The strategy of averaging the quantities over the scenarios with non-binding

capacity constraints performs well since the echelon stock levels are not ideal when the quantities are reduced by the capacity constraints. Preliminary experiments (omitted here) showed that the latter strategy performs better than using the maximum quantity, or averaging over all scenarios.

When applying the policy, for each period t with setup (i.e., $Y_{it} = 1$), the ordered quantity Q_{it} of each item i is computed to bring the echelon stocks E_{it} to the replenishment level S_{it} (i.e., $Q_{it} = S_{it} - E_{it}$ if $Y_{it} = 1$, and 0 otherwise). However, the quantities must respect the production capacities and the flow conservation constraints (i.e., the stocks of components must be large enough to produce the item). If a quantity violates the flow conservation constraint, it is reduced to the largest feasible quantity. If a resource capacity is violated, the quantities of all items processed with the resource are reduced by the same percentage. The Electronic Companion provides the algorithm with the detailed steps to infer the parameters and to use the S-Policy.

Two nearest scenario policies are also investigated. In each period t , the nearest scenario policy implements the decisions Q_{it}^ω associated with the optimization scenario ω with the demands from periods 1 to t nearest to the actual demands. A variant of this policy follows a path in the scenario tree. In each period t , this variant implements the decisions associated with the branch of the current path with demands in period t nearest to the actual demands. As preliminary experiments (not presented here) showed that these policies perform poorly, they were not further considered.

Proposition 5. *For the special case of the static-dynamic MMCLP with a single item and infinite capacity, there exists an optimal S-policy.*

Proof. See the Electronic Companion. □

The proposed S-policy was designed by extending the optimal rule for the special case of the MMCLP with a single item and infinite capacity (as stated in Proposition 5), thanks to the echelon stock computation. While the S-policy is optimal for this special case, it cannot be applied directly to the multi-echelon capacitated problem. Indeed, if the optimal replenishment quantities violate the capacity (or if they lead to negative levels of components inventory), a decision must be made to split the replenishment among the items. Therefore, the repairing procedure reduces the production quantities to satisfy these constraints. However, the experimental results show that the S-policy (which is a policy commonly used in practice) performs very well on the uncapacitated MMCLP but not on the capacitated MMCLP.

4.4 Rolling Horizon Framework

In practice, production planning tools are often used in a rolling horizon framework (Venkataraman 1996, Meistering and Stadtler 2017). Using the proposed methods in a rolling horizon framework leads to heuristics for the dynamic-dynamic decision framework. In the rolling horizon framework, the plan is optimized in period 0 by considering the first \mathcal{H} periods, and the decisions of period 0 are implemented. Then, the demands of period 0 are revealed, and the backlogs and inventories are observed. Considering this information, the plan is re-optimized on the horizon 1 to $\mathcal{H} + 1$. This process continues until the last period.

5 Comparison Methods and Simulation Framework

In this section, we describe the methods considered to benchmark the proposed approaches (Section 5.1) and the simulation framework considered to evaluate these approaches (Section 5.2).

5.1 Classical MRP Approaches

The proposed stochastic optimization approaches are compared with classical MRP approaches, namely, the deterministic mathematical model, and lot-sizing rules (lot-for-lot, economic order quantity, economic order period, Silver-Meal) equipped with safety stocks. Two approaches are considered to include safety stocks. As suggested in the literature (e.g., Zhao et al. 2001), the first approach assumes that the safety stocks are computed for end items only at the MPS level. The second approach computes the safety stocks with the method introduced in Graves and Willems (2008) for a base stock policy in multi-echelon supply chains with non-stationary demand. The Electronic Companion gives the detailed computation steps for these classical approaches.

5.2 Evaluation Methodology

5.2.1 Evaluation Framework

To estimate the expected total costs associated with the use of a method, a simulation is performed over 5,000 scenarios. These evaluation scenarios are different from the scenarios used for optimization, but they are sampled from the same distributions.

The simulation is performed independently on each scenario ω , and it results in the implementation of a solution s with setups Y_s and quantities Q_s . The solution s is directly available in the

static-static decision framework, whereas the simulation of the static-dynamic decisions plans the production in each period based on the latest information on the demand (as explained in the next section). The cost of s is computed using the deterministic model with scenario ω , where the setup and quantity variables are respectively fixed to Y_s and Q_s .

5.2.2 Re-planning Procedure

This section explains the re-planning methodology for the static-dynamic decision framework.

The solution of an instance \mathcal{P} of MMCLP gives the setups Y_{it} for each item i and period t , as well as the production quantity Q_{i0} in period 0. Given an evaluation scenario ω , the quantities $Q_{i\tau}$ to produce in periods $\tau > 0$ are decided sequentially by taking into account the information on the demands in periods $0, \dots, \tau - 1$. More precisely, $Q_{i\tau}$ is computed from the instance \mathcal{P}_τ^ω which differs from \mathcal{P} as follows: (1) The planning horizon becomes τ, \dots, T ; (2) The setups Y_{it} are given for t in τ, \dots, T ; (3) The initial inventory $I_{i\tau-1}^\omega$ and the backlog levels $B_{i\tau-1}^\omega$ are computed based the previously decided quantities $Q_{i0}, Q_{i1}^\omega, \dots, Q_{i\tau-1}^\omega$ and the demands $D_{i0}^\omega, \dots, D_{i\tau-1}^\omega$.

The considered methods require some adjustments to solve the modified instances \mathcal{P}_τ^ω . The *lot-sizing rules* set the quantity to 0 if there is no setup ($Q_{it} = 0$ if $Y_{it} = 0$), but they are applied as described in Section 5.1 otherwise ($Y_{it} = 1$). In the two-stage, multi-stage and deterministic models, the setup variables are set to the given values, thus a linear program is solved to determine the remaining continuous variables. To speed up the evaluation, the linear programs LP_τ associated with each instance \mathcal{P}_τ^ω are adjusted (and not completely re-built) to each evaluation scenario ω . More precisely, the values of the initial inventory levels (possibly negative) are updated, as well as the value of M_{it} to account for the production of the eventual backlogs. LP_τ is solved with the barrier method because our preliminary experiments (not presented here) showed that barrier performs faster than other linear programming solvers implemented in CPLEX for this problem. The *S-Policy* requires no modification, the echelon stock is computed based on the initial state, and the value Q_{i0}^ω is computed as described in Section 4.3.

The representation of the stochastic demand does not require any adjustment for the modified instances \mathcal{P}_τ^ω . More precisely, the lot-sizing rules and the deterministic models consider the average demand with safety stocks, whereas the two-stage and multi-stage models use samples of scenarios. Note that the multi-stage model considers scenario trees with structure $[N_1, \dots, N_{T-\tau}]$, where

$[N_1, \dots, N_T]$ refers to the structure of the optimization tree.

6 Numerical Experiments

The numerical experiments have two main objectives: (1) to evaluate the benefits of stochastic optimization for MRP systems when compared to classical safety stock approaches; (2) to investigate the performance of the proposed solution approaches for MRP under demand uncertainty.

Regarding objective (1), the results show that stochastic optimization models yield significantly lower costs than solving the deterministic model (or simple lot-sizing rules) with safety stocks. In addition, stochastic optimization yields solutions that result in lower volatility in the overall production cost. More precisely, adjusting the decisions when new information is available reduces the overall cost when planning with stochastic optimization, whereas updating the decisions increases the costs when planning with deterministic models or simple lot-sizing rules.

Regarding objective (2), the results show that 50 scenarios sampled with RQMC or QMC lead to good approximations of the two-stage model. Consequently, solving the two-stage model does not require a large computational effort. On the contrary, a good approximation of the multi-stage model requires 3,200 scenarios sampled with RQMC. Therefore, the multi-stage model is harder to solve. However, the two heuristics proposed in this work appear to be very efficient.

Table 3 summarizes the approaches evaluated in the experiments and their notations. The methods were implemented with Python and CPLEX 12.7, and run on an Intel(R) Xeon(R) X5675 3.07GHz processor. The code is available at <https://github.com/StochasticLotSizing/StochasticMRP>.

Section 6.1 introduces the considered instances. The performance of the sampling techniques and optimization approaches for stochastic optimization (objective 1) is shown in Section 6.2. Then, Section 6.3 (resp. 6.4) presents the simulation results for the static-static and static-dynamic (resp. dynamic-dynamic) decision framework that compares the proposed methods with the classical deterministic model with safety stocks (objective 2). Finally, Section 6.5 investigates the limits of the proposed fix-and-optimize approach on large supply chain planning instances.

Notation	Methods	Section
Average	The deterministic model without safety stocks	5.1
SS	The deterministic model with safety stocks	5.1
LL	The Lot-for-lot rule	5.1
EOQ	The Economic Order Quantity approach	5.1
SM	The Silver Mill approach	5.1
MPS	The safety stocks computed at the Master Planning level	5.1
GS	The safety stocks computed as in Graves and Willems (2008)l	5.1
2-stage	The two-stage model	3.2
M-stage	The multi-stage model	3.3
2-stage-H	The two-stage heuristic	4.2
Fix-&-Opt	The fix-and-optimize approach	4.2
S-Policy	The S policy	4.3
Q-Policy	The policy with fixed ordered quantity inferred from the two-stage model	3.2
CMC	The Crude Monte Carlo scenario sampling method	4.1
QMC	The Quasi Monte Carlo scenario sampling method	4.1
RQMC	The Randomized Quasi Monte Carlo scenario sampling method	4.1

Table 3: Summary of the notations used in the numerical experiments section.

6.1 MRP Instances Generation

We generate two test beds. First, the main experiments are performed with a test bed derived from the series A of Tempelmeier and Derstroff (1996). As these instances were designed for the deterministic MRP with zero lead times, they are extended (as explained in the Electronic Companion) to include the demand’s probability distributions and lead times. Three sets of instances are considered: (1) 1,026 *classical instances* are generated with full factorial design from the parameters given in Table 4; (2) 48 *instances with small distribution support* are generated with an assembly BOM (see Figure 2), a binomial distribution, and with full factorial design from the parameters (other than BOM and distribution) indicated in Table 4. (3) 20 *instances with a large planning horizon* are generated with lead times randomly chosen in $[0, 3]$, a time horizon of $\hat{T} + 10$ periods, and various values for the other parameters. Second, to evaluate the performance of the fix-and-optimize approach for large-scale supply chain planning problems, we generate a second test bed from the data in Willems (2008) as explained in the Electronic Companion.



Figure 2: Considered BOM.

Parameters	Values
BOM	General; Assembly (as depicted in Figure 2)
Resource structure	The items at the same echelon share the same resource
Resource utilization	Uncapacitated; 90%; 50%
Time between orders	1; 4
Distribution	Lumpy; Slow Moving; or Non-Stationary with known demand rate in $\{0.25; 0.5; 0.75\}$ and coefficient of variation in $\{0.1; 0.4; 0.7\}$
Lead time	L_1 (all items have a lead time of 1 period) L_2 (the lead times are equal to 1 for components, and 0 for end items)
Cost structures (ratio backlog/inventory costs)	2; 4
Echelon holding costs	Constant; large added value at last steps

Table 4: Values of the studied parameters.

6.2 Effectiveness of the Sampling and Solution Approaches

This section evaluates the impact of the scenario sampling techniques and of the number of scenarios on the two-stage and multi-stage models (Section 6.2.1). These experiments are performed with a subset (of size 20) of the classical instances with various structures. Then, the solutions obtained with samples of scenarios are compared with the true optimal solutions for the instances with small distribution supports, where the set of possible demand values is small enough to generate the full scenario set (Section 6.2.3). Finally, we evaluate with the subset (of size 20) of the classical instances the computation times and the performance of the stochastic models and heuristics for the static-dynamic decision framework (Section 6.2.2).

The performance of the methods is measured by the percentage gap (denoted GAP) between the expected total cost obtained with the method and the expected total cost obtained with the best performing method (best among the static-static and static-dynamic framework) on the considered instances. Unless otherwise stated, this measure is used in all the experiments.

6.2.1 Effectiveness of the Sampling Methods

Figure 3a shows the GAP of the two-stage model in the static-static environment for different sample sizes (10, 25, 50, 100, 200, and 500) and for the CMC and RQMC scenario generation methods. Figure 3a shows that a good approximation requires a sufficiently large scenario sample. For instance, a good approximation with CMC requires 200 scenarios. In addition, RQMC outperforms CMC (especially with small scenario samples), whereas the performance of QMC and RQMC are similar. Therefore, advanced scenario sampling techniques (such as RQMC) reduce the number of required scenarios. Indeed, 50 scenarios sampled with RQMC or QMC leads to good approxima-




Figure 3: Approximation quality with CMC, QMC, or RQMC and various numbers of scenarios.

tions, which are comparable to the solution obtained with 200 scenarios sampled with CMC.

Regarding the multi-stage model, a first set of experiments (omitted here) showed that the scenario tree structure does not have a strong impact on the approximation quality, but scenario trees with a large number of branches at the early stages and a reasonable number of branches at the last stages perform better when the model is re-solved in each period (i.e., the S-policy is not used). However, when the S-policy is used, the experiments show that a balanced scenario tree structure is more favorable than a scenario tree with a large number of scenarios in early stages. Figure 3b shows the *GAP* of the multi-stage model for various numbers (1600, 3200, 6400, 12800, and 25600) of scenarios sampled with CMC, QMC, and RQMC, and the computational performance is provided in the Electronic Companion. Figure 3b shows that RQMC leads to better solutions than CMC, especially when few scenarios are considered. In fact, 3,200 scenarios sampled with RQMC generally lead to a good approximation of the stochastic process, since considering a larger set of scenarios does not further reduce the *GAP*. Finally, QMC samples the same demands in all the nodes at the same level of the scenario tree. This leads to a poor representation of the stochastic demands compared to RQMC. In the rest of the experiments the multi-stage model is solved with 6,400 scenarios sampled with RQMC, since it is the largest number of scenarios which allows CPLEX to solve the entire problem to optimality within 10 hours (see the Electronic Companion). Even though the preliminary experiments consider instances with a relatively short time horizon, the multi-stage model can solve problems with a large time horizon by using a scenario tree with one branch at each node for the last periods in conjunction with a rolling horizon framework.

6.2.2 Numerical Results on Instances with a Perfect Set of Scenarios

The Electronic Companion reports the results of experiments on instances with small distribution supports, where the scenario set can be enumerated and the size is manageable. These experiments compare the sampling approximations with the true optimal solutions, and they show that RQMC scenario sampling yields near-optimal solutions. In addition, these results allow computing the expected value of perfect information (EVPI) which is equal to 55.96% and 49.77% for the static-static and static-dynamic decision frameworks, respectively.

6.2.3 Performance Comparisons of the Stochastic Approaches

The Electronic Companion compares the performance (in terms of solution quality and computation times) of the multi-stage model, the two-stage heuristic, and the fix-and-optimize heuristic in the static-dynamic environment. The results show that fix-and-optimize is efficient with a *GAP* of 0.30% versus 0.11% for the multi-stage model, whereas the multi-stage model is significantly slower to solve (241 seconds versus 3010 seconds on average). The two-stage heuristic leads to slightly larger costs, with a *GAP* of 0.73%, but it requires only 10.67 seconds to solve on average. In addition, the Electronic Companion reports the results obtained by five different runs of the methods. These results show that the methods are robust since different runs give similar results.

6.3 Effectiveness of Stochastic Models for MRP Under Demand Uncertainty

This subsection validates the effectiveness of stochastic optimization in MRP under demand uncertainty when compared to classical approaches. The results for the uncapacitated case are reported first to analyze the performance of the lot-sizing rules. Next, we report the results for the capacitated (more general) case, and we study the impact of the stochastic optimization approaches on MRP instances with different characteristics. According to the results presented in Section 6.2, the following solution approaches are considered: the two-stage model (denoted by *2-stage*) and the two-stage heuristic (denoted by *2-stage-H*) with 500 scenarios sampled with *RQMC*; the multi-stage model (denoted by *M-stage*) and the fix-and-optimize heuristic (denoted by *Fix-&-Opt*) using a scenario tree with structure [50, 8, 4, 4] (resulting in 6,400 scenarios); the *S-policy* determined by the multi-stage model using a scenario tree structure [10, 10, 8, 8] (resulting in 6,400 scenarios); the deterministic model with the average demands scenario (denoted *Average*), as well as the deter-

ministic model with safety stocks computed at the MPS level (denoted *SS-MPS*), and safety stocks computed using the guaranteed service time model (denoted *SS-GS*); the lot-sizing rules with safety stocks computed at the MPS level (*LL-MPS*, *EOQ-MPS*, *EOP-MPS*, *SM-MPS*), and safety stocks computed with the guaranteed service time model (*LL-GS*, *EOQ-GS*, *EOP-GS*, *SM-GS*). However, *SS-GS* is only considered in the static-dynamic framework because the safety stock of components cannot be transformed into end items in the static-static framework. Also, the lot-sizing rules are only considered for uncapacitated case, since preliminary experiments showed that they perform poorly on the capacitated case.

6.3.1 Uncapacitated Production Planning

Tables 5 and 6 report the numerical results on uncapacitated instances in the static-static environment. Table 5 gives the CPU time required to solve the problem at period 0, the *GAP*, the expected cost per component (setup, inventory, backlog, lost-sale, and production), and Table 6 reports some KPIs on the resulting production plan, namely, the proportion of demand delivered on time (fulfillment rate), backlogged, and in lost sales, as well as the average number of setups, and the expected number of periods covered by each lot. The 2-stage approach leads to the lowest costs (with a *GAP* of 3.1%), followed by *SS-MPS* (11.3%), and *SS-GS* (12.2%). Among the considered lot-sizing rules, *SM-MPS* and *SM-GS* perform the best, with a *GAP* of 18.9% and 19.2%, respectively.

Tables 7 and 8 report the results for the static-dynamic environment in the same format as Tables 5 and 6. The results show that Average, *LL-GS*, and *LL-MPS* lead to higher total costs in the static-dynamic environment than in the static-static. For instance, Average has a *GAP* of 49.6% in the static-static environment, versus 59.0% in static-dynamic. However, the safety stocks attenuate this problem, for instance, *SS-MPS* performs better in static-dynamic (8.9%) than in static-static (11.3%). Stochastic optimization approaches are superior in such circumstances, as 2-stage-H (0.6%) outperforms 2-stage (3.1%). Tables 6 and 8 show that these approaches yield better results with fewer setups (thus, more coverage) because producing larger lots helps to protect against demand uncertainty. This last remark shows the importance of the integration of lot sizes and safety stocks computations in MRP systems.

	CPU Time (sec.)	<i>GAP</i> (%)	Setup (\$)	Inventory (\$)	Backlog (\$)	Lost sales (\$)	Production (\$)
Average	0.0	49.6	7,151.6	3,528.2	1,856.1	9,707.5	3,370.3
SS-MPS	0.2	11.3	7,839.4	4,992.7	1,188.6	3,088.9	4,055.2
LL-MPS	0.2	58.3	23,761.8	2,565.3	614.2	2,093.5	4,033.4
EOQ-MPS	0.2	59.6	13,916.0	11,610.3	101.5	525.3	4,895.4
EOP-MPS	0.3	22.8	11,051.9	7,613.8	260.6	196.6	4,387.8
SM-MPS	0.5	18.9	12,523.7	5,621.6	307.7	196.6	4,387.8
LL-GS	3.7	60.1	23,775.1	2,315.0	642.8	2,554.2	3,952.3
EOQ-GS	3.7	59.3	13,766.1	11,454.8	103.0	700.6	4,822.3
EOP-GS	3.8	22.1	11,052.2	7,050.1	280.4	673.3	4,288.4
SM-GS	3.9	19.2	12,457.0	5,239.7	336.5	673.3	4,287.4
2-stage	7.7	3.1	7,077.2	5,071.0	1,807.7	1,048.8	3,983.2

Table 5: Results for the 352 uncapacitated instances in static-static environments.

	Fulfillment (%)	Backlog (%)	Lost sales (%)	No. setup	Coverage
Average	82.13	10.69	7.18	24.59	2.51
SS-MPS	89.53	8.72	1.75	25.39	2.70
LL-MPS	92.97	5.23	1.80	39.47	1.26
EOQ-MPS	97.70	1.73	0.57	30.11	2.61
EOP-MPS	97.46	2.35	0.19	21.48	3.03
SM-MPS	96.37	3.43	0.20	29.88	2.04
LL-GS	92.36	5.25	2.39	39.47	1.21
EOQ-GS	97.31	1.67	1.01	29.88	2.58
EOP-GS	96.99	2.45	0.56	21.47	2.88
SM-GS	95.70	3.70	0.60	29.80	1.95
2-stage	86.83	12.18	0.99	23.16	3.12

Table 6: Results for the 352 uncapacitated instances in static-static environments.

6.3.2 Capacitated Production Planning

This section compares the performance of the methods on all classical instances. Table 9 (resp. 10) reports the average *GAPs* for each method on each instance type for the static-static (resp. static-dynamic) environment, as well as the CPU time required to solve the instance in period 0. As for the uncapacitated case, Average performs better in the static-static environment than in static-dynamic. Moreover, these results show that stochastic optimization yield significant savings in the operating costs of production plans in a capacitated environment, since 2-stage significantly outperforms deterministic methods with safety stocks in static-static with an average *GAP* of 2.62% versus 8.36% for SS-MPS, as well as in static-dynamic with an average 0.63% for 2-stage-H, versus 6.21% and 6.80% for SS-MPS and SS-GS respectively.

The static-dynamic use of M-stage, 2-stage-H, Fix-&-Opt, Average, SS-MPS, and SS-GS is cumbersome because a linear program must be solved in each period. On the contrary, the static-

	CPU Time (sec.)	<i>GAP</i> (%)	Setup (\$)	Inventory (\$)	Backlog (\$)	Lost sales (\$)	Production (\$)
Average	0.0	59.0	7,151.6	3,455.9	1,950.3	10,879.0	3,304.1
SS-GS	10.6	10.0	7,779.0	4,574.1	1,255.0	3,554.5	3,904.1
SS-MPS	0.1	8.9	7,843.9	4,723.7	1,208.1	3,242.2	3,890.8
LL-MPS	0.2	64.0	23,765.8	2,291.3	691.9	3,132.2	3,833.6
EOQ-MPS	0.2	58.2	13,923.9	11,388.5	103.6	599.0	4,778.9
EOP-MPS	0.3	20.2	11,055.6	7,188.4	290.8	452.8	4,169.1
SM-MPS	0.5	16.9	12,528.2	5,223.4	344.2	469.0	4,132.8
LL-GS	3.7	64.0	23,775.1	2,182.0	700.4	3,142.7	3,820.3
EOQ-GS	3.7	57.6	13,766.1	11,280.9	106.8	725.0	4,758.2
EOP-GS	3.8	21.3	11,052.2	6,851.3	308.3	803.5	4,178.3
SM-GS	3.9	18.5	12,457.0	5,048.9	369.9	841.2	4,143.7
2-stage-H	7.7	0.6	7,077.2	4,831.8	1,831.7	1,044.6	3,891.9
S-Policy	310.2	0.6	7,322.5	4,866.0	1,719.0	857.3	3,893.4
Fix-&-Opt	232.0	0.2	7,077.2	4,988.1	1,797.4	840.8	3,924.4
M-stage	841.7	0.1	7,323.8	4,838.4	1,772.3	757.6	3,913.3

Table 7: Results for the 352 uncapacitated instances in static-dynamic environments.

	Fulfillment (%)	Backlog (%)	Lost sales (%)	No. setup	Coverage
Average	80.58	11.49	7.93	24.59	2.50
SS-GS	88.35	9.55	2.10	25.23	2.61
SS-MPS	89.00	9.13	1.87	25.45	2.64
LL-MPS	91.61	5.75	2.64	39.51	1.20
EOQ-MPS	97.48	1.72	0.81	30.10	2.55
EOP-MPS	97.03	2.55	0.41	21.51	2.93
SM-MPS	95.58	3.96	0.46	29.93	1.95
LL-GS	91.34	5.78	2.88	39.47	1.19
EOQ-GS	97.23	1.69	1.08	29.88	2.55
EOP-GS	96.73	2.61	0.66	21.47	2.85
SM-GS	95.16	4.10	0.74	29.80	1.91
2-stage-H	86.71	12.33	0.96	23.16	3.08
S-Policy	87.31	11.87	0.82	23.74	3.03
Fix-&-Opt	87.09	12.18	0.74	23.16	3.17
M-stage	87.33	12.00	0.67	23.89	3.04

Table 8: Results for the 352 uncapacitated instances in static-dynamic environments.

dynamic use of the S-policy does not require additional computations. Among the real-time execution methods, the S-policy performs the best with a *GAP* of 1.78%. However, investing more computation power reduces the costs, since M-stage leads to a *GAP* of 0.07% in the static-dynamic decision framework.

The gap between stochastic optimization methods and classical methods is exacerbated when:

- (1) The demand uncertainty is large (that is, the instances with lumpy demands, slow moving items, low rate of known demand). For instance, the *GAP* of SS-MPS increases from 4.57% to 7.70%, when the rate of known demand decreases from 75% to 25%.
- (2) The times required to transform the components into end items are short (that is, the instances with lead times of

type L2). (3) The value of the items increases significantly at each production step (that is, the instances with large echelon costs). In addition, M-stage significantly outperforms 2-stage-H in terms of solution quality in these situations. For instance, the *GAPs* of M-stage and 2-stage-H are respectively 0.05% and 1.90% for lumpy demands, versus 0.02% and 0.26% for non-stationary demands. These experimental findings complement the structural analysis presented in Section 3.4. Although Proposition 3 concerns a specific case of MMCLP, the experiments show that the M-stage significantly outperforms 2-stage-H for the generic MMCLP with short lead times. Similarly, the additional benefit of the M-stage (compared with 2-stage-H) is more pronounced when the difference between the holding costs of end items and components is large. Such situations are typically encountered for products with high added value in the supply chain.

Finally, all the methods (except the S-policy) have lower *GAPs* when capacity is tight. This is not surprising since the capacities constrain the solution. For instance, if the capacity is tight, the best production plan would be to produce as much as possible. However, S-policy performs poorly under tight capacities because of the greedy procedure used to repair capacity violations. Typically, when capacity is violated, the repair procedure reduces the quantity in equal proportion for all the items, whereas prioritizing some items might be preferable.

The Electronic Companion provides an analysis of the plans created with the considered approaches, which shows that the stocks of components are larger in the static-dynamic than in the static-static environment. Indeed, as stated in Propositions 1 and 2, the static-dynamic framework allocates the excess stock at the component level (with lower holding costs) when demand is low. Since 2-stage-H overlooks the possibility to react to low demand by keeping a stock of components, it overestimates the holding costs. Therefore, 2-stage-H creates smaller lots than M-stage and Fix-&-Opt.

The numerical results reported in Section 6.3 show that the use of stochastic optimization in MRP systems significantly reduces the operating costs. On the one hand, unlike the methods that determine lot sizes and safety stocks separately, the stochastic MRP models which explicitly take into account demand uncertainty in complex lot-sizing decisions can produce cost-effective production plans to hedge against demand uncertainty. On the other hand, the multi-stage stochastic model can account for the dynamic decision framework in a multi-echelon supply chain, which allows the production planner to proactively plan and react to the observed demand.

		Average	SS-GS	SS-MPS	2-stage
Distribution	NonStationary	30.53	6.45	6.34	1.39
	Lumpy	79.33	22.96	22.09	9.77
	SlowMoving	59.67	13.19	12.78	6.52
NonStationary Rate of known	25%	35.70	8.05	7.70	1.95
	50%	32.35	6.77	6.74	1.41
	75%	23.54	4.54	4.57	0.81
NonStationary Coeff. Variation	0.1	29.57	5.74	5.67	1.77
	0.4	29.73	5.79	5.81	1.13
	0.7	32.30	7.82	7.53	1.27
BOM	Assembly	32.06	6.15	5.84	1.50
	General	43.17	10.98	10.87	3.74
Utilization	50%	19.77	5.15	5.81	1.72
	90%	43.48	8.36	7.96	3.08
	Uncapacitated	49.60	12.18	11.30	3.05
TBO	1	63.37	9.43	9.24	4.70
	3	11.87	7.70	7.47	0.54
Lead times	L1	34.71	6.39	6.67	1.05
	L2	40.52	10.74	10.04	4.18
Echelon cost	Normal	22.92	6.71	6.90	1.52
	Large	52.31	10.42	9.81	3.72
Cost structure	2	22.89	6.46	6.91	2.35
	4	52.34	10.67	9.80	2.88
All instances		37.62	8.57	8.36	2.62
CPU Time		0.10	10.59	0.17	8.95

Table 9: GAPs of the methods in the static-static decision framework on the 1056 classical instances.

6.4 Rolling Horizon Simulation

In practice, production planning tools are often used in a rolling horizon framework (Chand et al. 2002). This section presents a rolling horizon simulation and studies the impact of the considered planning horizon. The rolling horizon simulation is performed in a dynamic-dynamic decision framework on instances with large time horizon ($\mathcal{T} = \hat{T} + 10$ periods), and with different planning horizons $\mathcal{H} < \mathcal{T}$. In period 0, the simulation optimizes the plan by considering the first \mathcal{H} periods, and it implements the resulting decisions for period 0. Then, the revelation of the real demands in period 0 allows one to observe the backlogs and inventories. With this information, the simulation re-optimizes the plan on the horizon 1 to $\mathcal{H} + 1$. This process continues until the last period. To keep the duration of the evaluation process reasonable, the simulation uses 100 scenarios on each instance. Indeed, with a time horizon of 10 periods, the evaluation over 100 scenarios requires to re-solve the problem 1,000 times per instance.

		Average	SS-GS	SS-MPS	S-Policy	2-stage	Fix-&-Opt	M-stage
Distribution	NonStationary	38.82	5.31	5.00	0.93	0.26	0.03	0.02
	Lumpy	103.42	18.26	15.44	5.18	1.90	0.96	0.05
	SlowMoving	69.66	8.84	7.79	6.03	2.69	1.78	0.55
NonStationary Rate of known	25%	45.16	6.41	5.97	1.22	0.34	0.04	0.02
	50%	41.11	5.65	5.37	0.96	0.26	0.03	0.01
	75%	30.20	3.86	3.67	0.60	0.16	0.02	0.02
NonStationary Coeff. Variation	0.1	38.10	4.37	3.90	0.81	0.31	0.01	0.00
	0.4	38.69	4.87	4.71	1.05	0.32	0.04	0.01
	0.7	39.68	6.67	6.40	0.91	0.14	0.04	0.03
BOM	Assembly	44.50	5.34	4.50	0.96	0.50	0.20	0.12
	General	50.50	8.27	7.92	2.60	0.75	0.35	0.01
Utilization	50%	29.19	3.98	4.08	3.83	0.56	0.27	0.01
	90%	54.27	6.49	5.61	0.90	0.73	0.33	0.10
	Uncapacitated	59.03	9.94	8.92	0.60	0.59	0.22	0.09
TBO	1	79.84	6.31	5.17	2.34	1.05	0.44	0.07
	3	15.16	7.29	7.25	1.22	0.20	0.11	0.07
Lead times	L1	44.12	5.80	5.61	1.16	0.29	0.17	0.08
	L2	50.88	7.81	6.80	2.39	0.96	0.38	0.06
Echelon cost	Normal	30.50	5.94	5.33	2.04	0.26	0.15	0.08
	Large	64.50	7.67	7.08	1.51	1.00	0.40	0.05
Cost structure	2	29.79	4.87	4.71	1.56	0.61	0.26	0.05
	4	65.21	8.74	7.70	1.99	0.64	0.29	0.09
All instances		47.50	6.80	6.21	1.78	0.63	0.27	0.07
CPU Time		0.10	10.59	0.17	1064.03	8.95	234.16	3240.36

Table 10: GAPS of the methods in the static-dynamic decision framework on the 1056 classical instances

The following methods are considered in the rolling horizon framework: *Average*, *SS-MPS*, and *SS-GS* with a planning horizon of $(\hat{T} + 5)$ periods; *2-stage-H1*, *2-stage-H3*, and *2-stage-H5* denote respectively the two-stage heuristic with time-horizons of $(\hat{T} + 1)$, $(\hat{T} + 3)$, and $(\hat{T} + 5)$ periods; *Fix-&-Opt-H1*, *Fix-&-Opt-H3*, and *Fix-&-Opt-H5* denote the fix-and-optimize heuristic with time-horizons of $(\hat{T} + 1)$, $(\hat{T} + 3)$, and $(\hat{T} + 5)$ periods. Following the results presented in Section 6.2, the fix-and-optimize heuristic uses a scenario tree structure [50, 8, 4, 4] for the first four levels, whereas the nodes of subsequent levels have one branch. We also evaluate two policies which do not require to solve a mathematical model in each period. In *S-Policy-H6*, the S-Policy is inferred (as explained in Section 4.3) from the solution of the fix-and-optimize heuristic with a planning horizon of $(\hat{T} + 6)$ periods. The model is first solved on the planning horizon $0 \dots (\hat{T} + 6)$, and the resulting S-policy is executed in periods $0 \dots 2$. Then, the model is solved with the planning horizon $3 \dots (\hat{T} + 9)$, and the S-Policy is executed in periods $3 \dots 5$. This process continues until the end of the horizon.

In *Q-Policy-H6*, the two-stage model is solved over a planning horizon of $(\hat{T} + 6)$ periods, and the resulting production quantities are implemented for the first 3 periods. Then, the two-stage model is re-solved by shifting the planning horizon by 3 periods (similarly to S-Policy-H6). In addition, to compare the static-static decision framework and the rolling horizon approach, *Average*, *SS-MPS*, and *2-stage* are simulated in the static-static decision framework (i.e., the plan remains fixed). However, *2-stage* is solved with 100 scenarios only at each iteration, because the two-stage model with 500 scenarios per iteration requires too much memory on these instances. Though the multi-stage model could be applied in a rolling-horizon framework, the evaluation over a large number of scenarios is not considered because it requires a significantly higher computation time.

Tables 11 and 12 report the detailed costs and the KPIs (in the same format as Table 5 and 6). The results show that the benefit of stochastic optimization in MRP is even more pronounced in the dynamic-dynamic environment than in the static-static and static-dynamic environments, as the resulting plans lead to significantly lower costs than classical methods. The methods based on safety stocks such as SS-MPS generally yield better results in the static-static decision framework than in the rolling horizon framework. On the contrary, the results show that the stochastic approaches perform better in a rolling horizon framework than in the static-static framework. For instance, 2-stage-H5 in the rolling horizon framework outperforms 2-stage in the static-static framework. As expected, using a large planning horizon leads to lower costs. For instance, the *GAP* of the fix-and-optimize method decreases from 13.7% to 1.5% when the planning horizon increases from $(\hat{T} + 1)$ to $(\hat{T} + 5)$. Finally, Q-Policy-H6 outperforms S-Policy with *GAPs* of 3.6% versus 5.2%. As mentioned earlier, S-Policy does not perform well in presence of capacity. However, on the subset of 12 uncapacitated instances, S-Policy outperforms Q-Policy (3.9% versus 4.4%).

6.5 Numerical Tests on Large-Scale Supply Chain Instances

This section evaluates the performance and the limits of Fix-&-Opt for large-scale supply chain management problems with the instances derived from the real-world instances of Willems (2008). These instances include both the production (MRP) and distribution (DRP) systems, and they are also used in recent academic works (e.g., Kumar and Aouam 2019). We run Fix-&-Opt with 100 scenarios in the first step, and a scenario tree with structure [50, 8, 4, 4] in the second step. We limit the CPU time of step 1 to one hour to avoid excessive computing requirements. Table

Framework	Method	GAP (%)	Setup (\$)	Inventory (\$)	Backlog (\$)	Lost sales (\$)	Production (\$)
Static-static for entire horizon	Average	24.9	16,861.7	11,241.2	4,520.2	9,868.0	6,432.3
	SS-MPS	10.0	19,902.2	12,949.8	3,274.3	2,212.4	7,207.9
	2-stage	6.6	17,382.1	12,527.6	4,915.3	1,781.7	7,175.2
Rolling horizon	Average H5	43.8	20,422.0	8,382.4	7,639.3	15,626.8	6,115.2
	SS-GS H5	18.7	20,108.9	11,165.3	5,821.7	3,714.7	6,811.8
	SS-MPS H5	24.2	20,094.3	10,453.6	5,887.6	6,650.4	6,620.2
	2-stage H1	17.2	26,315.4	6,733.0	8,976.4	3,197.6	6,662.0
	Fix-&-Opt H1	13.7	25,956.9	7,452.5	7,758.7	1,923.1	6,756.4
	2-stage H3	5.1	19,898.8	9,871.1	5,707.8	2,388.4	6,700.1
	Fix-&-Opt H3	3.2	19,403.4	10,627.1	5,431.8	1,298.7	6,819.1
	2-stage H5	3.3	19,049.8	9,510.0	5,779.0	2,638.8	6,701.5
	Fix-&-Opt H5	1.5	18,810.9	10,572.8	4,997.4	1,365.5	6,823.5
	S-Policy H6	5.2	18,523.0	11,241.4	4,855.2	1,913.3	6,769.3
Q-Policy H6	3.6	18,774.0	10,684.3	4,994.1	1,926.2	6,816.5	

Table 11: Results of the rolling horizon simulation on instances with large horizon

Framework	Method	Fulfillment (%)	Backlog (%)	Lost sales (%)	No. setup	Coverage
Static-static for entire horizon	Average	74.6	16.2	9.2	53.7	2.7
	SS-MPS	80.9	16.9	2.2	57.7	2.7
	2-stage	79.3	18.7	2.1	42.3	3.4
Rolling horizon	Average H5	64.0	26.7	9.3	50.6	2.2
	SS-GS H5	74.3	23.7	2.0	53.2	2.6
	SS-MPS H5	72.0	24.0	4.0	52.3	2.5
	2-stage H1	71.6	25.9	2.5	60.5	2.0
	Fix-&-Opt H1	74.1	24.2	1.7	57.8	2.2
	2-stage H3	74.5	23.1	2.4	46.8	2.7
	Fix-&-Opt H3	77.2	21.2	1.6	44.4	2.9
	2-stage H5	75.7	21.6	2.7	44.0	2.8
	Fix-&-Opt H5	78.4	20.1	1.6	42.7	3.0
	S-Policy H6	77.7	19.9	2.4	44.6	2.9
Q-Policy H6	78.7	19.3	2.0	43.9	2.9	

Table 12: Results of the rolling horizon simulation on instances with large horizon

13 reports the number of items $|\mathcal{I}|$ and periods $|\mathcal{T}|$ in each instance, as well as the CPU time, the integrality gap (for step 1 only), and the number of variables and constraints in each step. Table 13 shows that the proposed approach performs well since it can solve realistic instances with up to 58 items and 17 periods. For a few instances, CPLEX cannot solve to optimality the two-stage model corresponding to the first step. However, the gap determined by the approach is relatively small (less than 4% in all the instances).

SC	I	T	Step 1				Step 2			
			CPU (s)	CPLEX Gap	No. Var.	No. Const.	CPU (s)	No. Var.	No. Const.	
1	8	7	2.06	0.00	1,777	899	2.15	99,135	69,671	
2	13	9	2.25	0.00	4,185	2,115	11.52	57,025	8,434	
3	17	11	434.05	0.00	3,776	1,949	21.04	69,023	8,652	
4	22	9	48.27	0.00	8,044	4,078	17.82	113,240	2,397	
5	27	11	3,600.02	0.04	9,704	4,960	36.72	150,137	3,466	
6	28	9	3,600.01	0.01	11,134	5,638	17.50	162,372	5,057	
7	38	9	14.61	0.00	1,071	677	16.66	406,801	355,047	
8	40	17	3,617.53	0.00	3,698	2,262	377.48	212,873	85,713	
9	49	11	3,600.07	0.01	27,987	14,110	102.52	384,769	27,332	
10	58	5	1.05	0.00	1,232	751	9.07	567,705	464,947	

Table 13: Computation time for Fix-&-Opt heuristic on large instances

7 Future Research Directions

This paper shows that stochastic optimization approaches can significantly reduce the production costs for companies relying on MRP systems in the static and dynamic decision frameworks. The fix-and-optimize heuristic is proven to be efficient for the more complex static-dynamic framework and could be used to solve realistic size MRP instances. The scalability of the approaches can be further enhanced in the following ways. In the first step of the fix-and-optimize heuristic, approaches based on a shifting window (Toledo et al. 2015) can find good quality solutions for the two-stage model in reasonable time. In the second step, scenario-wise or stage-wise decomposition methods can circumvent the issue related to the exponential growth in the number of scenarios.

Several stage-wise decomposition approaches exist for multi-stage stochastic optimization, such as nested Benders decomposition, Benders decomposition with block separation (Golari et al. 2017), stochastic decomposition (Sen and Zhou 2014), among others. Stochastic dual dynamic programming (SDDP) is particularly appealing for the static-dynamic decision framework since it avoids the exponential growth of the scenario tree. Provided that the demand probability distributions are stage-wise independent, the considered problem can be decomposed into a series of decision problems (one for each decision stage). SDDP solves the problem stage by stage, and it approximates the future costs with a set of cuts. Besides, Rebennack (2016) proposes to combine the nested Benders decomposition method with SDDP to incorporate additional sources of uncertainty that are too complex to model with SDDP (e.g., stage-wise dependent random parameters).

Progressive hedging (PH) decomposes the problem per scenario. For instance, Huang and Zheng (2020) use PH to solve a preventive maintenance scheduling problem. To decompose the problem

per scenario, PH relaxes the non-anticipativity constraints, but it penalizes their violations with a penalty similar to a Lagrangian multiplier. PH can remove the memory issue encountered in the second step of fix-and-optimize since the scenarios are solved one by one. While PH cannot solve the original MMCLP since the setup variables are integer (thus the problem is not convex), it can find the optimal solution to the second step of the fix-and-optimize approach.

In addition, the study of the complex dynamic-dynamic decision framework is an interesting avenue for future research. For instance, recent developments of the SDDP approach can handle integer recourse variables (Zou et al. 2018, 2019). A different approach is to use decision rules to transform the multi-stage model into a two-stage one. However, the design of such rules is complex for the MMCLP because of the capacity constraint. Such rules could also be combined with Lagrangian decomposition (Daryalal et al. 2020).

8 Managerial insights and conclusions

Manufacturers are widely using MRP systems to plan their production. The most recent MRP systems use mathematical programs to compute the lot sizes under the assumption of a deterministic demand, and separately computed safety stocks to hedge against uncertainty. To jointly optimize the decisions on the lot sizes and safety stocks, this paper investigates the use of stochastic optimization in MRP systems. More precisely, we provide a two-stage and a multi-stage formulation for the multi-echelon multi-item capacitated lot-sizing problem encountered in MRP systems. The two-stage model assumes that production quantities are fixed for the entire horizon. On the contrary, the multi-stage model represents the static-dynamic decision framework, where the production quantities are sequentially determined in each period.

The experimental results show that computing safety stocks separately from the lot sizes results in solutions that are far from being optimal. Therefore, our results suggest that manufacturing companies should consider adopting stochastic optimization approaches that integrate the computations of lot sizes and safety stocks in a single framework, which potentially leads to significant cost savings. Contrarily to the analytical computation of safety stocks, the scenario based stochastic optimization approach is highly flexible as it can include various manufacturing constraints, such as planned substitutions and component routing, and it does not require restrictive assumptions on the probability distribution. In addition, the scenario based approach can accommodate other

sources of uncertainty including lead times and process duration.

While the two-stage model corresponding to the static-static environment can be directly implemented in MRP systems, the multi-stage model requires heavy computational time. Thus, to efficiently determine MRP production plans, the two heuristics proposed in this work can be employed in conjunction with an order-up-to-level policy derived from the solution of the stochastic MRP model to make real-time recourse decisions during the static-dynamic execution. These approaches are efficient and easy to implement since they are designed to work with standard MILP solvers.

Acknowledgement

The authors thank the Departmental Editor, the Senior Editor, and five anonymous referees for their comments and suggestions on an earlier draft of the paper. The authors are also thankful to Maxime Cohen for helpful feedback on the revised manuscript. This research was supported by the Institut de Valorisation des Données (IVADO) and MITACS.

References

- Absi, N., S. Kedad-Sidhoum, S. Dauzère-Pérès. 2011. Uncapacitated lot-sizing problem with production time windows, early productions, backlogs and lost sales. *International Journal of Production Research* **49**(9) 2551–2566.
- Aloulou, M. A., A. Dolgui, M. Y. Kovalyov. 2014. A bibliography of non-deterministic lot-sizing models. *International Journal of Production Research* **52**(8) 2293–2310.
- Bai, X., J. S. Davis, J. J. Kanet, S. Cantrell, J. W. Patterson. 2002. Schedule instability, service level and cost in a material requirements planning system. *International Journal of Production Research* **40**(7) 1725–1758.
- Benton, W.C. 1991. Safety stock and service levels in periodic review inventory systems. *Journal of the Operational Research Society* **42**(12) 1087–1095.
- Billington, P. J., J. O. McClain, L J. Thomas. 1983. Mathematical programming approaches to capacity-constrained mrp systems: Review, formulation and problem reduction. *Management Science* **29**(10) 1126–1141.
- Bitran, G. R., H. H. Yanasse. 1982. Computational complexity of the capacitated lot size problem. *Management Science* **28**(10) 1174–1186.

- Blackburn, J. D., D. H. Kropp, R. A. Millen. 1985. MRP system nervousness: Causes and cures. *Engineering Costs and Production Economics* **9**(1-3) 141–146.
- Bookbinder, J. H., J.-Y. Tan. 1988. Strategies for the probabilistic lot-sizing problem with service-level constraints. *Management Science* **34**(9) 1096–1108.
- Boulaksil, Y. 2016. Safety stock placement in supply chains with demand forecast updates. *Operations Research Perspectives* **3**(1) 27–31.
- Brandimarte, P. 2006. Multi-item capacitated lot-sizing with demand uncertainty. *International Journal of Production Research* **44**(15) 2997–3022.
- Carlson, R. C., C. A. Yano. 1986. Safety stocks in MRP-systems with emergency setups for components. *Management Science* **32**(4) 403–412.
- Cecere, L. 2015. Does better forecasting improve inventory? Why I don't think so anymore. <https://www.forbes.com/sites/loracecere/2015/11/29/does-better-forecasting-improve-inventory-why-i-dont-think-so-anymore/amp/> (19/09/2020).
- Chand, S., V. N. Hsu, S. Sethi. 2002. Forecast, solution, and rolling horizons in operations management problems: A classified bibliography. *Manufacturing & Service Operations Management* **4**(1) 25–43.
- Chevreur, L., M. Hu, S. Gandhi. 2018. Why supply chains must pivot. <https://sloanreview.mit.edu/article/why-supply-chains-must-pivot/> (19/09/2020).
- Clark, A. R., V. A. Armentano. 1995. A heuristic for a resource-capacitated multi-stage lot-sizing problem with lead times. *Journal of the Operational Research Society* **46**(10) 1208–1222.
- Daryalal, M., M. Bodur, J. R. Luedtke. 2020. Lagrangian dual decision rules for multistage stochastic mixed integer programming. *arXiv preprint* arXiv:2001.00761.
- Dolgui, A., C. Prodhon. 2007. Supply planning under uncertainties in MRP environments: A state of the art. *Annual Reviews in Control* **31**(2) 269–279.
- Dupačová, J., N. Gröwe-Kuska, W. Römisch. 2003. Scenario reduction in stochastic programming. *Mathematical Programming* **95**(3) 493–511.
- Enns, S. T. 2002. MRP performance effects due to forecast bias and demand uncertainty. *European Journal of Operational Research* **138**(1) 87–102.
- Golari, M., N. Fan, T. Jin. 2017. Multistage stochastic optimization for production-inventory planning with intermittent renewable energy. *Production and Operations Management* **26**(3) 409–425.
- Graves, S. C., T. Schoenmeyr. 2016. Strategic safety-stock placement in supply chains with capacity constraints. *Manufacturing & Service Operations Management* **18**(3) 445–460.
- Graves, S. C., S. P. Willems. 2008. Strategic inventory placement in supply chains: Nonstationary demand.

- Manufacturing & Service Operations Management* **10**(2) 278–287.
- Grubbström, R. W., Z. Wang. 2003. A stochastic model of multi-level/multi-stage capacity-constrained production–inventory systems. *International Journal of Production Economics* **81** 483–494.
- Guide, V. D. R., R. Srivastava. 2000. A review of techniques for buffering against uncertainty with MRP systems. *Production Planning & Control* **11**(3) 223–233.
- Ho, C.-J., Tim C. Ireland. 1998. Correlating MRP system nervousness with forecast errors. *International Journal of Production Research* **36**(8) 2285–2299.
- Huang, Z., Q. P. Zheng. 2020. A multistage stochastic programming approach for preventive maintenance scheduling of GENCOs with natural gas contract. *European Journal of Operational Research* **287**(3) 1036–1051.
- Inderfurth, K. 2009. How to protect against demand and yield risks in MRP systems. *International Journal of Production Economics* **121**(2) 474–481.
- Inderfurth, K., S. Minner. 1998. Safety stocks in multi-stage inventory systems under different service measures. *European Journal of Operational Research* **106**(1) 57–73.
- Kadipasaoglu, S. N., V. Sridharan. 1995. Alternative approaches for reducing schedule instability in multi-stage manufacturing under demand uncertainty. *Journal of Operations Management* **13**(3) 193–211.
- Kaut, M., S. W. Wallace. 2007. Evaluation of scenario-generation methods for stochastic programming. *Pacific Journal of Optimization* **3**(2) 257–271.
- Kumar, K., T. Aouam. 2019. Extending the strategic safety stock placement model to consider tactical production smoothing. *European Journal of Operational Research* **279**(2) 429–448.
- Lagodimos, A. G., E. J. Anderson. 1993. Optimal positioning of safety stocks in MRP. *The International Journal of Production Research* **31**(8) 1797–1813.
- L’Ecuyer, P., D. Munger. 2016. Algorithm 958: Lattice builder: A general software tool for constructing rank-1 lattice rules. *ACM Transactions on Mathematical Software (TOMS)* **42**(2) 15.
- Lin, P.-C., R. Uzsoy. 2016. Chance-constrained formulations in rolling horizon production planning: an experimental study. *International Journal of Production Research* **54**(13) 3927–3942.
- Meistering, M., H. Stadtler. 2017. Stabilized-cycle strategy for capacitated lot sizing with multiple products: Fill-rate constraints in rolling schedules. *Production and Operations Management* **26**(12) 2247–2265.
- Mordor Intelligence. 2020. Enterprise resource planning market-growth, trends, and forecast (2020-2025). [https://www.orbisresearch.com/reports/index/enterprise-resource-planning-market-growth-trends-and-forecast-2020-2025\(19/09/2020\)](https://www.orbisresearch.com/reports/index/enterprise-resource-planning-market-growth-trends-and-forecast-2020-2025(19/09/2020)).
- Rebennack, S. 2016. Combining sampling-based and scenario-based nested Benders decomposition methods:

- application to stochastic dual dynamic programming. *Mathematical Programming* **156**(1-2) 343–389.
- Sali, M., V. Giard. 2015. Monitoring the production of a supply chain with a revisited MRP approach. *Production Planning & Control* **26**(10) 769–785.
- Sen, S., Z. Zhou. 2014. Multistage stochastic decomposition: a bridge between stochastic programming and approximate dynamic programming. *SIAM Journal on Optimization* **24**(1) 127–153.
- Tempelmeier, H. 2013. Stochastic lot sizing problems. *Handbook of Stochastic Models and Analysis of Manufacturing System Operations*. Springer, 313–344.
- Tempelmeier, H., M. Derstroff. 1996. A Lagrangean-based heuristic for dynamic multilevel multiitem constrained lotsizing with setup times. *Management Science* **42**(5) 738–757.
- Toledo, C. F. M., M. da Silva Arantes, M. Y. B. Hossomi, P. M. França, K. Akartunali. 2015. A relax-and-fix with fix-and-optimize heuristic applied to multi-level lot-sizing problems. *Journal of Heuristics* **21**(5) 687–717.
- Venkataraman, R. 1996. Frequency of replanning in a rolling horizon master production schedule for a process industry environment: A case study. *Production and Operations Management* **5**(3) 255–265.
- Willems, S. P. 2008. Real-world multiechelon supply chains used for inventory optimization. *Manufacturing & Service Operations Management* **10**(1) 19–23.
- Zahorik, A., L. J. Thomas, W. W. Trigeiro. 1984. Network programming models for production scheduling in multi-stage, multi-item capacitated systems. *Management Science* **30**(3) 308–325.
- Zhang, Y., Z.-J. Max Shen, S. Song. 2016. Distributionally robust optimization of two-stage lot-sizing problems. *Production and Operations Management* **25**(12) 2116–2131.
- Zhao, X., F. Lai, T. S. Lee. 2001. Evaluation of safety stock methods in multilevel material requirements planning (mrp) systems. *Production Planning & Control* **12**(8) 794–803.
- Zhao, X., T.S. Lee. 1993. Freezing the master production schedule for material requirements planning systems under demand uncertainty. *Journal of Operations Management* **11**(2) 185–205.
- Zijm, H., G.-J. Van Houtum. 1994. On multi-stage production/inventory systems under stochastic demand. *International Journal of Production Economics* **35**(1-3) 391–400.
- Zou, J., S. Ahmed, X. A. Sun. 2018. Multistage stochastic unit commitment using stochastic dual dynamic integer programming. *IEEE Transactions on Power Systems* **34**(3) 1814–1823.
- Zou, J., S. Ahmed, X. A. Sun. 2019. Stochastic dual dynamic integer programming. *Mathematical Programming* **175**(1-2) 461–502.