European Journal of Operational Research 302 (2022) 221–229

Contents lists available at ScienceDirect

European Journal of Operational Research

journal homepage: www.elsevier.com/locate/ejor

Production, Manufacturing, Transportation and Logistics

Dynamic lot sizing with stochastic demand timing

Kerem Akartunalı[a,*], Stéphane Dauzère-Pérès[b,c]

[a] Department of Management Science, University of Strathclyde, 130 Rottenrow, Glasgow G4 0GE, UK
[b] Mines Saint-Etienne, Univ. Clermont Auvergne, CNRS, UMR 6158 LIMOS, CMP, Department of Manufacturing Sciences and Logistics, Gardanne F-13541, France
[c] Department of Accounting and Operations Management, BI Norwegian Business School, 0484 Oslo, Norway

ARTICLE INFO

Article history:
Received 1 April 2021
Accepted 17 December 2021
Available online 24 December 2021

Keywords:
Production
Lot Sizing
Dynamic Programming
Stochastic Demand Timing

ABSTRACT

In this paper, a novel way of modeling uncertainty on demand in the single-item dynamic lot sizing problem is proposed and studied. The uncertainty is not related to the demand quantity, but rather to the demand timing, i.e., the demand fully occurs in a single period of a given time interval with a given probability and no partial delivery is allowed. The problem is first motivated and modeled. Our modeling naturally correlates uncertain demands in different periods contrary to most of the literature in lot sizing. Dynamic programs are then proposed for the general case of multiple demands with stochastic demand timing and for several special cases. We also show that the most general case where the backlog cost depends both on the time period and the stochastic demand is NP-hard.

1. Introduction

This paper tackles a single-item dynamic lot sizing problem, i.e., quantities to be produced or replenished on a finite planning horizon discretized in periods must be determined to satisfy time-varying demands. The total cost, which combines fixed setup costs and variable inventory and production costs, is to be minimized. Because we consider uncertainty, backlog costs associated with delaying the satisfaction of uncertain demands in a period are also included in the total cost.

Handling the uncertainty of parameters in planning problems is a challenging task. A straightforward way (though naive and often costly) is to use buffers to cover for random events, such as safety stocks to cover for larger demands than expected in production and inventory planning. Another classical and more complex approach is to explicitly consider the probability distributions of the stochastic parameters and to minimize the total expected cost. Most of the literature in lot sizing investigates deterministic problems, see, e.g., Brahimi, Absi, Dauzère-Pérès, & Nordli (2017) for an excellent review of single-item problems, Doostmohammadi & Akartunalı (2018) for a recent theoretical overview on complex multi-item problems, and Meistering & Stadtler (2017) for a recent overview of problems with rolling schedules. Recent surveys on stochastic lot sizing can be found in Tempelmeier (2013) and Aloulou, Dolgui, & Kovalyov (2014). In stochastic lot sizing problems, the total expected (discounted or not) cost is often minimized (or respectively, profit maximized). A set of scenarios can be used to model the problem effectively, as in Guan & Miller (2008) to design a polynomial time algorithm for the most simplistic case of a single item, or as in Golari, Fan, & Jin (2017) to develop a sophisticated decomposition approach for real-world problems with multi-stage decision making. An alternative to expected cost is to use service levels, modeled through chance constraints, as in Tempelmeier (2007).

In their survey, Brahimi et al. (2017) show (see Table 2) that the vast majority of the research literature in single-item stochastic dynamic lot sizing investigate stochastic demands with a particular focus on volumes. Stochastic demands have been taken in account in some lot sizing models with pricing decisions (Federgruen & Heching, 1999; Thomas, 1974), and stochastic costs and yield have also been studied in combination with stochastic demands, e.g., in Huang & Ahmed (2010). On the other hand, stochastic lead times have very rarely been considered. In their work, Huang & Küçükyavuz (2008) address the single-item problem with stochastic lead times and propose a dynamic programming algorithm that is polynomial in the size of the scenario tree to solve the problem, which was later improved by Jiang & Guan (2011). It is also noteworthy to remark the approximation algorithms proposed by Levi & Shi (2013) for lot sizing problems with stochastic lead times.

We note that robust optimization approaches have also been used to handle demand uncertainties in lot sizing since the earlier works of Ben-Tal, Golany, Nemirovski, & Vial (2005) and

* Corresponding author
E-mail addresses: kerem.akartunali@strath.ac.uk (K. Akartunalı), dauzere-peres@emse.fr (S. Dauzère-Pérès).

https://doi.org/10.1016/j.ejor.2021.12.027

Bertsimas & Thiele (2006). The exact min-max decomposition approach of Bienstock & Özbay (2008) is further extended by Attila, Agra, Akartunalı, & Arulselvan (2021) to include uncertainties in returns in a remanufacturing setting. Though limited due to its static nature, Wei, Li, & Cai (2011) propose a robust LP formulation, and the general dynamic programming framework of Agra, Santos, Nace, & Poss (2016) is shown to work effectively in lot sizing problems. Distributionally robust optimization has also been shown to be an effective tool for two-stage decision making in practice (Zhang, Shen, & Son, 2016). However, even more noticeable than the stochastic lot sizing literature, the focus has been solely on uncertainties in the volume of demand. A recent review of the broad field of robust optimization can be found in Gorissen, Yanikoğlu, & den Hertog (2015).

Our problem setting significantly differs from previous studies, and partly answers one of the main criticisms associated with the modeling of stochastic demands in lot sizing problems, namely that uncertain demands in different periods are uncorrelated. Stochastic demands are usually considered as independent random variables, which is unrealistic in many practical settings, where an increase of the demand in one period is often associated with a decrease of the demand in another period. In our modeling, we consider that the demands are deterministic in terms of volumes but that their timing might be stochastic. More precisely, a given demand quantity might occur in a window of multiple periods, with a given probability to fully occur in each of these periods. Hence, stochastic demands are naturally correlated due to the fact that the total demand on the horizon is deterministic while the occurrences of the demands are stochastic.

Stochastic demand timing can be observed in a range of practical settings. In particular, this happens when a client company sends orders for a product to a supplier company, when the client company's product inventory level is empty. The order to the supplier is fixed, typically related to the inventory capacity of the customer. Hence, the supplier company knows very well the quantity that will be either picked up by or delivered to a customer, but is not able to know exactly when, although an interval of several days is known. This is particularly noticeable in operational or tactical production and inventory planning over several weeks with periods of a day, where demand and order quantities are well established, and is a typical context in process industries, which satisfy the demands of other industries. For example, this case is observed for non-mixable cement products that can be stored (see, e.g., Christiansen et al. (2011)) or calcium carbonate slurry products (see, e.g., Dauzère-Pérès et al. (2007)), where it is known that a vessel, a train, or a truck will arrive in an interval of several days to be completely filled.

Hence, order management is an interesting context, where stochastic demand timing is relevant. When a company has a list of potential customers' orders, predicted from historical data, with known quantities and time windows in which they should occur with their corresponding probabilities, solving the problem studied in this paper will provide the most efficient plan to answer these orders. Significant potential losses due to future orders can thus be estimated, and necessary actions to avoid these losses can be taken.

More generally, manufacturers may be able to very efficiently forecast the demand quantity from a given retailer in the upcoming weeks, which must be delivered in a single period, but will have more difficulty to forecast the exact timing of such demand. Stochastic demand timing as considered in this paper models this general context.

We remark that the time windows considered in this paper are different from the demand (or delivery) time windows introduced in Lee, Çetinkaya, & Wagelmans (2001), and from the production time windows introduced in Dauzère-Pérès, Brahimi, Najid, & Nordli (2002). Demand time windows specifically correspond to grace periods, in which demands can be delivered without incurring holding or backlog costs, whereas this is not the case in our problem. On the other hand, production time windows require that each demand quantity be produced within a specified time window, while there are no such constraints in our problem. However, combining stochastic demand timing with demand or production time windows could be a potential area for future investigation.

The problem is first formalized and analyzed in Section 2. Then, in Section 3, we study the simplest possible case of stochastic demand timing, i.e., a single interval, in order to facilitate the discussion of more complex cases later. A polynomial dynamic program is proposed and some properties are introduced. The general case of stochastic demand timing with multiple intervals is then analyzed in Section 4, where some additional results are presented and a dynamic program of exponential complexity in the worst case is proposed. We show that this dynamic program is polynomial in the realistic case, where production costs are time independent and the ratio between inventory and backlog costs is time independent, and also in the case where probability distributions are convex. We then study in Section 5 two special cases with further assumptions on the intervals of stochastic demand timing: Firstly, in Section 5.1, we assume the intervals to be non-overlapping, and secondly, in Section 5.2, we assume that a dominance order exists between the intervals. Both cases are shown to be polynomial. In Section 6, we extend the results of the previous sections to the case where the backlog cost depends not only on the period but also on the quantity of stochastic demand, and establish its complexity to be $\mathcal{NP}$-hard. Finally, we conclude with key remarks and future research directions in Section 7.

2. Problem modeling

Let us consider the single-item uncapacitated dynamic lot sizing problem with a planning horizon of T periods in the classical deterministic sense, as follows:

$$\min \sum_{t=1}^{T} f_t y_t + \sum_{t=1}^{T} h_t s_t + \sum_{t=1}^{T} c_t x_t \tag{1}$$

$$\text{s.t. } x_t + s_{t-1} - s_t = D_t \qquad t = 1, \ldots, T \tag{2}$$

$$x_t \le M_t y_t \qquad t = 1, \ldots, T \tag{3}$$

$$y_t \in \{0, 1\}; x_t \ge 0; s_t \ge 0 \qquad t = 1, \ldots, T \tag{4}$$

For any period t, variables x_t and s_t represent production and inventory quantities, respectively, and binary y_t variables indicate whether a production setup takes place or not. The objective (1) is to find a minimum cost production plan, where the total cost consists of fixed setup costs f_t (charged only if production is strictly positive, i.e., $y_t = 1$), per unit inventory holding costs h_t, and per unit production costs c_t, respectively, for all periods in the horizon. We also assume all cost parameters to be strictly positive, i.e., no "free lunch". The flow balance constraints (2) ensure on-time satisfaction of demand D_t, whereas the relationship between production and setup variables is set by (3), where M_t is an upper bound on x_t, e.g., $M_t = \sum_{\ell=t}^{T} D_\ell$. Finally, the integrality and non-negativity constraints are provided by (4). Let us recall that this problem has a complexity of $O(T \log T)$, see, e.g., Wagelmans, van Hoesel, & Kolen (1992).

In addition to the deterministic demands D_t, $\forall t \in [1, T]$, that need to be satisfied on time, we simultaneously consider stochastic demand timing as follows. Let $[l_i, u_i] \subset [1, T]$ be an interval, indexed by i, where it is certain that a demand of d^i will fully occur, i.e., at once, in one period, with a probability of $p_t^i \ge 0$ for each period $t \in [l_i, u_i]$ and such that $\sum_{t=l_i}^{u_i} p_t^i = 1$. Note that $p_t^i = 0$

for $t \leq l_i - 1$ and $t \geq u_i + 1$. Let $\mathcal{I}$ be the set of such intervals with stochastic demand timing in the planning horizon and, for ease of notation, let $|\mathcal{I}| = n$.

In this paper, we make the following realistic assumptions:

- No backlog is allowed for deterministic demands and, accordingly, no backlog is allowed for any stochastic demand quantity d^i after period u_i. Note that, however, stochastic demand quantity d^i may be satisfied with inventory carried from before l_i, while backlogging is allowed within the interval $[l_i, u_i]$ with a variable backlog cost b_t. In Section 6, the more general case where the variable backlog cost b_t^i also depends on d^i is discussed.
- Partial delivery of any stochastic demand quantity d^i is not allowed, i.e., d^i products must be delivered to the customer in one and only one period. Hence, each stochastic demand timing can be seen as a separate order, and the backlog cost is counted until d^i is fully satisfied. Note that the problem is easy to solve if partial delivery is allowed, as one can simply solve in that case a classical lot sizing problem with demand $p_t^i d^i$ in period t.
- As it is usually the case and w.l.o.g., backlog is more costly than inventory, i.e., $b_t > h_t \; \forall t$.

For any period $t \in [l_i, u_i]$, the expected stochastic demand quantity to satisfy is $p_t^i d^i$. As this stochastic demand quantity cannot be produced after u_i, we note that, for any $t \leq u_i$ and per unit produced, the expected inventory is $\sum_{l=t+1}^{u_i} p_l^i$ (if one unit of product has already been produced) and the expected backlog is $\sum_{l=l_i}^{t-1} p_l^i$ (if one unit of product has not been produced yet). Hence, the expected holding and backlog cost for producing one unit of product to satisfy d^i in period t is denoted by $EC_i(t)$, which can be defined as follows for any $t \leq u_i$:

$$EC_i(t) = \sum_{l=t}^{u_i} h_l \sum_{k=l+1}^{u_i} p_k^i + \sum_{l=l_i}^{t-1} b_l \sum_{k=l_i}^{l} p_k^i \tag{5}$$

Note that the first and second terms of (5) correspond to the expected holding and backlogging costs, respectively. Also, note that the first term is equal to 0 for $t = u_i$, and the second term is equal to 0 for $t \leq l_i$. Next, we present a numerical example to illustrate the problem.

Example 1. Consider a problem with five periods and two stochastic demand timing intervals, i.e., $T = 5$, $n = 2$. For the sake of simplicity, let the cost parameters be time independent, and let $h_t = 1.5$, $b_t = 6$, $f_t = 25$ and $c_t = 8$, $t = 1, \ldots, 5$. The remaining parameter values are given as follows:

t	1	2	3	4	5	
D_t	4	0	10	6	9	
p_t^1	0.45	0.35	0.2	0	0	$d^1 = 7$, $[l_1, u_1] = [1, 3]$
p_t^2	0	0	0.3	0.7	0	$d^2 = 5$, $[l_2, u_2] = [3, 4]$

We first note that in period 5, at most 9 units will be produced, i.e., the deterministic demand of period 5, and no stochastic demand. On the other hand, in the first three periods, d^1 and/or d^2 can be produced, while in period 4, d^2 can be produced, in addition to any deterministic demand that is produced. To illustrate (5), we provide the following detailed calculations for the cases of producing in period t when $l_i < t < u_i$, $t < l_i$ and $t = u_i$:

$EC_1(2) = h_2 p_3^1 + b_1 p_1^1 = 1.5\times0.2 + 6\times0.45 = 3$

$EC_2(1) = h_1(p_3^2 + p_4^2) + h_2(p_3^2 + p_4^2) + h_3 p_4^2 = 1.5\times1 + 1.5\times1 + 1.5\times0.7 = 4.05$

$EC_1(3) = b_1 p_1^1 + b_2(p_1^1 + p_2^1) = 6\times0.45 + 6\times0.8 = 7.5$

Recall that these are unit costs for expected holding and backlogging costs. For example, producing one unit of d^1 in period 2 will incur an expected cost of 3, in addition to the unit production cost of 8 and fixed cost of 25. □

The inventory variable s_t is a stochastic variable since d^i is stochastic, and thus modeling our problem by extending the model (1)-(4) is not trivial. Hence, as it is common in lot sizing, we propose to formalize our problem with the variables in $[0, 1]$ z_{lt}, the fraction of the deterministic demand D_t produced in period $l \leq t$, and z_l^i, the fraction of the stochastic demand quantity d^i produced in period $l \leq u_i$. In order to illustrate the development of our model, we first reformulate the deterministic model (1)-(4) using the z_{lt} variables, which are linked to the original production variables as follows:

$$x_l = \sum_{t=l}^{T} z_{lt} D_t, \quad l = 1, \ldots, T. \tag{6}$$

Then, the deterministic model becomes:

$$\min \sum_{t=1}^{T} f_t y_t + \sum_{t=1}^{T} \sum_{l=1}^{t} (c_l + \sum_{k=l}^{t-1} h_k) z_{lt} D_t \tag{7}$$

$$\text{s.t.} \sum_{l=1}^{t} z_{lt} = 1 \qquad t = 1, \ldots, T \tag{8}$$

$$\sum_{t=l}^{T} z_{lt} D_t \leq M_l y_l \qquad l = 1, \ldots, T \tag{9}$$

$$y_t \in \{0, 1\} \qquad t = 1, \ldots, T \tag{10}$$

$$0 \leq z_{lt} \leq 1 \qquad t = 1, \ldots, T; \; l = 1, \ldots, t \tag{11}$$

We remark that the objective (1) is rewritten as (7) using (6) and the fact that inventory variables are no longer explicitly used. Constraints (8) ensure that the deterministic demands are satisfied in the horizon, and constraints (9) correspond to constraints (3) using (6).

Then, using z_l^i associated with the stochastic demand quantities, we next state the relationship between the original production variables and the new variables in a similar fashion to (6):

$$x_l = \sum_{t=l}^{T} z_{lt} D_t + \sum_{i \in \mathcal{I};\; l \leq u_i} z_l^i d^i, \quad l = 1, \ldots, T. \tag{12}$$

Our problem can then be modeled as follows:

$$\min \sum_{t=1}^{T} f_t y_t + \sum_{t=1}^{T} \sum_{l=1}^{t} (c_l + \sum_{k=l}^{t-1} h_k) z_{lt} D_t + \sum_{i \in \mathcal{I}} \sum_{l=1}^{u_i} (c_l + EC_i(l)) z_l^i d^i \tag{13}$$

$$\text{s.t.} (8), (11)$$

$$\sum_{l=1}^{u_i} z_l^i = 1 \qquad i \in \mathcal{I} \tag{14}$$

$$\sum_{t=l}^{T} z_{lt} D_t + \sum_{i \in \mathcal{I};\; l \leq u_i} z_l^i d^i \leq M_l y_l \qquad l = 1, \ldots, T \tag{15}$$

$$y_t \in \{0, 1\} \qquad t = 1, \ldots, T \tag{16}$$

$$0 \leq z_l^i \leq 1 \qquad i \in \mathcal{I}; \; l = 1, \ldots, u_i \tag{17}$$

In a similar fashion to (7), the objective (1) is rewritten as (13) using (12) and (5). Constraints (14) ensure that the stochastic demand quantities are produced in the horizon similar to constraints (8) for deterministic demands. Constraints (15) correspond to constraints (3) using (12).

Next, we remark the following result.

Proposition 1. $\arg\min EC_i(t) \in [l_i, u_i]$.

Proof. First, note that the first term of (5) is strictly decreasing over $[1, u_i]$ since $h_t > 0\ \forall t$, while the second term of (5) is strictly increasing over $[l_i, u_i]$ since $b_t > 0\ \forall t$. To prove that the minimum of $EC_i(t)$ is attained in $[l_i, u_i]$, it is sufficient to observe that the second term of (5) is 0 for $t \le l_i$ while the first term of (5) attains its lowest value over $[1, l_i]$ at $t = l_i$. □

In the remainder of the paper, and for the sake of simplicity, we use the notation t_i^* to indicate the period where the minimum of $EC_i(t)$ is attained, i.e. $t_i^* = \arg\min EC_i(t)$. In case of multiple periods attaining this minimum, we assume that t_i^* indicates the earliest such period. Finally, we note that the problem can be rewritten with only stochastic demand quantities by considering that $p_t^i = 1$ and $l_i = t = u_i$ for D_t.

3. Stochastic demand timing with a single interval

In this section, we assume that there is a single interval i with stochastic demand timing, with a demand quantity of d^i throughout the planning horizon. Because backlog on d^i is only allowed before u_i, d^i is either produced before or at l_i, i.e., no backlog cost is incurred, or between $l_i + 1$ and u_i, i.e., both inventory and backlog costs are incurred.

The following theorem states that there is an optimal solution in which d^i is not produced in multiple periods, and that if stochastic demand quantity is produced, it is not produced in isolation from deterministic demand, thus limiting the number of states in the dynamic program.

Theorem 1. *There is always an optimal solution, where demand d^i is produced in a single period $t \le u_i$. Moreover, $x_t \ge d^i + D_t$ holds when there are no speculative production costs, i.e., $c_t + \sum_{\ell=t}^{t'-1} h_\ell \ge c_{t'}$, $\forall t, t' \in [1, T]$ such that $t < t'$.*

Proof. Let us consider an optimal solution where d^i is produced in two periods t' and t'', i.e. the setups costs f_t and $f_{t'}$ are both incurred. It can be observed that the expected total cost can be reduced by producing d^i only in period t' if $c_{t'} + EC_i(t') \le c_{t''} + EC_i(t'')$ or only in period t'' if $c_{t''} + EC_i(t'') \le c_{t'} + EC_i(t')$. Next, assume that d^i is produced in period t in an optimal solution, and $x_t = d^i$. Then, if $D_t > 0$, production in a period $t' < t$ must include D_t. However, producing D_t in t would save $(c_{t'} + \sum_{\ell=t'}^{t-1} h_\ell - c_t)D_t$, where $c_{t'} + \sum_{\ell=t'}^{t-1} h_\ell - c_t \ge 0$ when production costs are not speculative. □

Theorem 1 is used in all the dynamic programs proposed in this paper by only considering solutions where the quantity d^i of each stochastic demand timing interval $[l_i, u_i]$ is produced in a single period. Moreover, as shown in the proof of Proposition 1, the minimum of $EC_i(t)$ over $[1, l_i]$ is reached at l_i. Hence, we consider in the dynamic programs that the decisions to produce d^i are made between l_i and u_i, even if d^i is produced between 1 and $l_i - 1$.

To solve the case with a single stochastic demand timing interval, the main change from the dynamic program of Wagner & Whitin (1958) is that two states are managed for $t \in [l_i + 1, u_i]$, depending on whether the decision of producing d^i has already been taken or not. Hence, we make the following state definition for each period t and indicator parameter $sd^i \in \{0, 1\}$ in order to derive the dynamic programming algorithm:

$G(t, sd^i)$: The value of the optimal solution for the horizon $[1, t-1]$, where $sd^i = 1$ indicates the case of demand d^i being already produced before t, and $sd^i = 0$ indicates the case of demand d^i not being produced yet.

By problem definition, $G(1, 0) = 0$ holds, and we also note that for $t \le l_i$, only $sd^i = 0$ is allowed, whereas for $t \ge u_i + 1$, only $sd^i = 1$ is allowed. Next, we define $c_{t'k} = c_{t'} + \sum_{l=t'}^{k-1} h_l$ to represent the cost of meeting one demand unit of period k by producing in period t', where $t' \le k$. The recursion is then formally defined as follows:

- For $t \le l_i$ (i.e., d^i is still not produced):

$$G(t, 0) = \min_{t' \le t-1} \left(G(t', 0) + f_{t'} + \sum_{k=t'}^{t-1} c_{t'k} D_k \right)$$

- For $l_i + 1 \le t \le u_i$ (i.e., d^i is produced or not):

$$G(t, 0) = \min_{t' \le t-1} \left(G(t', 0) + f_{t'} + \sum_{k=t'}^{t-1} c_{t'k} D_k \right)$$

$$G(t, 1) = \min_{t' \le t-1} \left(G(t', 0) + f_{t'} + \sum_{k=t'}^{t-1} c_{t'k} D_k + d^i (c_{t'} + EC_i(t')), \right.$$
$$\left. G(t', 1) + f_{t'} + \sum_{k=t'}^{t-1} c_{t'k} D_k \right)$$

- For $t \ge u_i + 1$ (i.e., d^i must be produced):

$$G(t, 1) = \min \left(\min_{t' \le u_i} \left(G(t', 0) + f_{t'} + \sum_{k=t'}^{t-1} c_{t'k} D_k + d^i (c_{t'} + EC_i(t')) \right), \right.$$
$$\left. \min_{t' \ge l_i+1} \left(G(t', 1) + f_{t'} + \sum_{k=t'}^{t-1} c_{t'k} D_k \right) \right).$$

The optimal cost will be given by $G(T + 1, 1)$. Following the $O(T logT)$ algorithms proposed in Wagelmans et al. (1992) for the case without backlogging cost, and in van Hoesel (1991), Aggarwal & Park (1993) and Federgruen & Tzur (1993) for the case with backlogging cost, it is straightforward to observe that our dynamic program can also be implemented with a complexity of $O(T logT)$. Note that it is also possible to first solve the problem in $O(T logT)$ with only the deterministic demands, and if there is already a setup in period t_i^*, then the stochastic demand quantity can be added. Finally, we remark that if the production costs are non-speculative, then the complexity reduces to $O(T)$, in line with previous results such as presented in Wagelmans et al. (1992).

4. General case of stochastic demand timing

First, we investigate the general case of stochastic demand timing, in order to propose a general purpose dynamic programming algorithm. As we will discuss later, this algorithm will be improved from a computational complexity perspective when more restricted but realistic special cases are considered.

Starting from the simplest case discussed in Section 3, the first obvious step for generalization is to consider multiple intervals with stochastic demand timing. Then, one can observe that such intervals may also have overlaps. Less obvious is a case when there is no particular order between such overlapping intervals, and therefore, we next define an essential property, in order to differentiate different cases of overlapping intervals.

Definition 1. Let d^i and d^j be two demands with stochastic timing. If $\sum_{l=1}^{t} p_l^i \ge \sum_{l=1}^{t} p_l^j\ \forall t \in [l_j, u_i]$, then we say that d^i **dominates** d^j.

Example 2. Consider a problem with five periods and three stochastic demand timing intervals, i.e., $T=5, n=3$. Assume we are given the following data for these intervals:

t	1	2	3	4	5	
p_t^1	0.1	0.5	0.4	0	0	$[l_1,u_1]=[1,3]$
p_t^2	0	0.3	0.2	0.5	0	$[l_2,u_2]=[2,4]$
p_t^3	0	0	0.6	0.2	0.2	$[l_3,u_3]=[3,5]$

Demand d^1 dominates d^2 since $0.1+0.5 \geq 0.3$ and $0.1+0.5+0.4 \geq 0.3+0.2$ both hold. On the other hand, neither d^2 nor d^3 dominate the other, since $0.3+0.2 \leq 0.6$ holds while $0.3+0.2+0.5 \geq 0.6+0.2$ is true. □

In this section, we look into the general case with multiple intervals of stochastic demand timing, where we do not have any dominance relationship between the overlapping intervals.

Let us also introduce the following definition, where we assume that $EC_i(t)=+\infty$ if $t \geq u_i+1$.

Definition 2. Let σ_i denote the sequence of length T for demand d^i in which periods are ranked in non-decreasing order of the production and expected unit holding and backlog cost $c_t + EC_i(t)$. More precisely, $\forall k=2,\ldots,T$, either i) $c_{\sigma_i(k)}+EC_i(\sigma_i(k)) > c_{\sigma_i(k-1)}+EC_i(\sigma_i(k-1))$ or ii) both $c_{\sigma_i(k)}+EC_i(\sigma_i(k)) = c_{\sigma_i(k-1)}+EC_i(\sigma_i(k-1))$ and $\sigma_i(k) > \sigma_i(k-1)$ hold.

Example 3. Using the first interval (i.e., $i=1$) from Example 2, suppose that $c_1+EC_1(1)=12$, $c_2+EC_1(2)=11$ and $c_3+EC_1(3)=15$ (note this is simply $+\infty$ for periods 4 and 5). Then, by a slight abuse of notation, our ordering vector is $\sigma_1=(2,1,3,4,5)$. □

Then, we propose the following result.

Theorem 2. *For two demands with stochastic timing d^i and d^j, if $\sigma_i=\sigma_j$, then there is an optimal solution in which d^i and d^j are produced in the same period.*

Proof. From Theorem 1, we know that there is an optimal solution in which d^i is produced in a single period t' and d^j is produced in a single period t''. If $\sigma_i=\sigma_j$ and $t'' \neq t'$ then, by definition of σ_i, the solution is only optimal if $c_{t'}+EC_i(t')=c_{t''}+EC_i(t'')$, otherwise the solution could be strictly improved by producing both demands d^i and d^j in t' if $c_{t'}+EC_i(t') < c_{t''}+EC_i(t'')$ or in t'' if $c_{t'}+EC_i(t') > c_{t''}+EC_i(t'')$. Finally, because $c_{t'}+EC_i(t')=c_{t''}+EC_i(t'')$, it is possible to change the solution and keep the same total cost by producing both demands d^i and d^j in t' or in t''. □

Theorem 2 implies that, for two stochastic demand timings such that $\sigma_i=\sigma_j$ and $u_i<u_j$, there is an optimal solution in which d_j is not produced between u_i+1 and u_j. Note also that there are $O(T!)$ possible different sequences of periods in σ_i.

4.1. Dynamic program for the general case

Let $(sd^1,\cdots,sd^n)$ be a vector of binary parameters, where sd^i is defined for each stochastic demand timing interval $i \in \mathcal{I}$ in the same fashion as in Section 3. Then, for the general dynamic program, we define $G(t,(sd^1,\cdots,sd^n))$, which indicates the value of the optimal solution for the horizon $[1,t-1]$ and the specific vector $(sd^1,\cdots,sd^n)$.

Note that a vector $(sd^1,\cdots,sd^n)$ is classified as *valid* at period t (or equivalently, $G(t,(sd^1,\cdots,sd^n))$ is valid) if:

- $sd^i=0$ for all $i \in \mathcal{I}$ such that $t \leq l_i$,
- $sd^i=0$ or $sd^i=1$ for all $i \in \mathcal{I}$ such that $t \in [l_i+1,u_i]$, and
- $sd^i=1$ for all $i \in \mathcal{I}$ such that $t \geq u_i+1$.

By definition, $G(1,(sd^1,\cdots,sd^n))=0$ holds, where $sd^i=0$, $\forall i \in \mathcal{I}$. Let $\mathcal{SD}(t)$ denote the set of valid vectors at period t. For each vector $(sd^1,\cdots,sd^n) \in \mathcal{SD}(t)$, the recursion for $G(t,(sd^1,\cdots,sd^n))$ is formally defined as follows:

$$G(t,(sd^1,\cdots,sd^n)) = \min_{\substack{t' \leq t-1,\\ (sd'^1,\cdots,sd'^n) \in \mathcal{SD}(t')}} \Bigg(G(t',(sd'^1,\cdots,sd'^n)) + f_{t'} + \sum_{k=t'}^{t-1} c_{t'k} D_k + \sum_{\substack{i \in \mathcal{I}:\\ sd^i - sd'^i = 1}} d^i (c_{t'} + EC_i(t')) \Bigg) \quad (18)$$

Note that Theorem 1 remains valid in this case, as one can extend these results by simply applying them to any interval $i \in \mathcal{I}$. Hence, the optimal cost for the full problem is given by $G(T+1,(sd^1,\cdots,sd^n))$, where $sd^i=1$, $\forall i \in \mathcal{I}$. We remark that, when $n=1$, i.e., there is a single interval, it is easy to observe that this general dynamic program exactly maps to the one described in Section 3: $G(t,sd^i)$ is reduced to a single stochastic demand timing while the validity arguments for sd^i remain (though now for a single interval), and the cost of producing d^i is only applied when sd^i value is changed from 0 to 1 in the new time period.

The complexity of the dynamic program is $O(T \max_{t \in [1,T]} |\mathcal{SD}(t)|)$. The value of $\max_{t \in [1,T]} |\mathcal{SD}(t)|$ is discussed in Lemma 1.

Lemma 1. *In the worst case,* $\max_{t \in [1,T]} |\mathcal{SD}(t)| \sim O(\min\{2^n, T!\})$

Proof. The worst case can be reached in two different ways:

1. If there exists $t \in [1,T]$ such that $t \in [l_i+1,u_i]$, $\forall i \in \mathcal{I}$, i.e., all n intervals intersect with each other at least in one period. This leads to $O(2^n)$ combinations.
2. Following Theorem 2, it is possible to combine demands with the same sequence σ_i in the same indicator sd^i in the dynamic programming algorithm. This leads to a maximum of $O(T!)$ combinations. This is essentially a preprocessing stage to the algorithm.

□

Therefore, the time complexity of the algorithm may be exponential in n and in T. However, as stochastic demand intervals will be short in most practical settings (no more than 4 or 5 periods), a small number of intervals should be overlapping in any period t, leading to small sets $\mathcal{SD}(t)$. For example, in case all demand intervals are different from each other (i.e., no two intervals have the same starting and ending points) and the number of periods of each interval is limited to m periods, then $\mathcal{SD}(t)=O(m^2)$. Moreover, if at most k intervals are overlapping in any period t, then the complexity of the dynamic program is $O(T2^k)$. Moreover, as we will see in Sections 4.2 and 4.3 for practical general cases, as well as in Section 5 for some relevant special cases, this time complexity can be effectively reduced to polynomial.

4.2. Time independent production costs and time independent ratio between inventory and backlog costs

An interesting case in practice appears when the ratio between the unit inventory and backlog costs in each period is time independent, i.e., $h_t=\alpha_t h$ and $b_t=\alpha_t b$ with $\alpha_t>0$ $\forall t$ (or, equivalently, $h_t/b_t=h/b$, $\forall t$). Moreover, we assume time independent production costs, i.e., $c_t=c$, $\forall t \in [1,T]$. Although this case is more restricted than the general case that does not specify cost functions or other key parameters of the problem, it is very common in practice, where hard to quantify backlog costs are often defined in terms of inventory holding costs. Moreover, its limitations are minimal, as there is no specification on how the actual cost levels would vary from one period to another, and time independent production costs are a common setting in the lot sizing literature. Because production costs are time independent, they can be ignored in the remainder of this section. Finally, as discussed in this section, this case can be solved in polynomial time. A more special case worth remarking is when the inventory and backlog costs are time independent, i.e., $\alpha_t = 1\ \forall t$.

First, we note the step change from t to $t+1$:

$$\Delta(t) = EC_i(t+1) - EC_i(t) = \left(\sum_{l=t+1}^{u_i} h_l \sum_{k=l+1}^{u_i} p_k^i + \sum_{l=l_i}^{t} b_l \sum_{k=l_i}^{l} p_k^i\right) - \left(\sum_{l=t}^{u_i} h_l \sum_{k=l+1}^{u_i} p_k^i + \sum_{l=l_i}^{t-1} b_l \sum_{k=l_i}^{l} p_k^i\right) = b_t \sum_{k=l_i}^{t} p_k^i - h_t \sum_{k=t+1}^{u_i} p_k^i$$

Because $\sum_{k=t+1}^{u_i} p_k^i = 1 - \sum_{k=l_i}^{t} p_k^i$, the expression above can be rewritten:

$$\Delta(t) = (h_t + b_t)\sum_{k=l_i}^{t} p_k^i - h_t \tag{19}$$

Note that (19) can also be used to show Proposition 1, since $\Delta(t) = -h_t < 0$ when $t \le l_i - 1$. For the case of time independent ratio, we can rewrite the expression (19) as follows:

$$\Delta(t) = \alpha_t\left((h+b)\sum_{k=l_i}^{t} p_k^i - h\right) \tag{20}$$

Theorem 3. *If the ratio between the inventory and backlog costs is time independent, i.e., $h_t = \alpha_t h$ and $b_t = \alpha_t b\ \forall t$, then $EC_i(t)$ is strictly decreasing until $t = t_i^*$ and strictly non-decreasing after $t = t_i^*$. Moreover, if inventory and backlog costs are time independent, i.e., $\alpha_t = 1\ \forall t$, then $EC_i(t)$ is convex.*

Proof. We first observe that, in (20), $\sum_{k=l_i}^{t} p_k^i$ is strictly increasing with t when $l_i \le t \le u_i$ (while being 0 when $t \le l_i - 1$, as noted earlier). Since $h + b > 0$, the value of $\Delta(t)$, starting from $-\alpha_t h < 0$ at $t = l_i - 1$, will also be strictly increasing. Hence, either (i) $t = t_i^* \le u_i - 1$ holds due to the first observation of $(h+b)\sum_{k=l_i}^{t} p_k^i \ge h$ at t, or (ii) $t_i^* = u_i$ holds if $EC_i(u_i) - EC_i(u_i - 1) < 0$. In case (i), note that $(h+b)\sum_{k=l_i}^{t} p_k^i = h$ is possible, and hence the function $EC_i(t)$ is strictly non-decreasing (rather than strictly increasing). This concludes the proof of the first claim.

Next, consider the case of $\alpha_t = 1$. Note that we can further simplify (20) by eliminating α_t. Then, we have:

$$EC_i(t+1) = EC_i(t) + (h+b)\sum_{k=l_i}^{t} p_k^i - h$$

$$EC_i(t+2) = EC_i(t) + (h+b)\sum_{k=l_i}^{t} p_k^i - h + (h+b)\sum_{k=l_i}^{t+1} p_k^i - h$$

where the second equation is simply the definition of $\Delta(t+1)$ with $EC_i(t+1)$ substituted using the first equation. Since $\sum_{k=l_i}^{t} p_k^i \le \sum_{k=l_i}^{t+1} p_k^i$, it is possible to observe that $EC_i(t+2) + EC_i(t) \ge 2EC_i(t+1)$. This concludes the convexity of $EC_i(t)$. □

The case of a convex $EC_i(t)$ function can be associated to the practical setting where, as one moves further away from $t = t_i^*$, not only the expected cost increases, but also the rate of the cost increases.

In line with the previous literature, we next define a regeneration interval $[t_1, t_2]$ as an interval of periods such that production takes place in periods t_1 and t_2 while no production occurs in periods t, $t_1 < t < t_2$. Then, we have the following result.

Proposition 2. *Given a regeneration interval $[t_1, t_2]$, let $\mathcal{I}_{t_1,t_2} = \{i \in \mathcal{I} : t_1 \le t_i^* \le t_2\}$. If the ratio between the inventory and backlog costs is time independent, and production costs are time independent, then in an optimal solution involving regeneration interval $[t_1, t_2]$, for every $i \in \mathcal{I}_{t_1,t_2}$, d^i will be produced either at t_1 or t_2.*

Proof. First, note that the production of d^i for any $i \in \mathcal{I}_{i_1,i_2}$ cannot take place in a period $t < t_1$ (or $t > t_2$), since $EC_i(t) \ge EC_i(t_1)$ (or $EC_i(t) \ge EC_i(t_2)$, respectively) due to Theorem 3 and the fact that production costs are time independent. Since production of d^i for any $i \in \mathcal{I}$ takes place in a single period in an optimal solution due to Theorem 1, and since, by definition, there is no production in any period t such that $t_1 < t < t_2$, d^i will be produced either in t_1 if $EC_i(t_1) \le EC_i(t_2)$, or in t_2 otherwise. □

Next, we discuss how to use this result to define a dynamic program of polynomial complexity particularly due to the significantly reduced number of linkages between states. First, we note that the number of valid states is reduced, since now a state is valid only if $sd^i = 0$ for all $i \in \mathcal{I}$ such that $t \le t_i^*$ (rather than $t \le l_i$). Next, in order to account for the regeneration intervals, we replace $\mathcal{SD}(t')$ with $\mathcal{SD}(t', t)$ in the recursion (18) of the dynamic program, where we define any valid $\mathcal{SD}(t', t)$ as follows:

- If $t' \le t_i^* \le t - 1$ and $EC_i(t') \le EC_i(t)$, then $sd^i = 1$ must hold at t,
- If $t' \le t_i^* \le t - 1$ and $EC_i(t') > EC_i(t)$, then $sd^i = 0$ must hold at t,
- If $t' \ge t_i^* + 1$, then $sd^i = 1$ must hold at t.

Note that the first case means that d^i must be produced at t' (since it is cheaper at t') whereas the second case means that d^i will be not produced at t'. In the third case, if $sd^i = 0$ holds at t', then d^i must be produced at t' since producing at t will be more expensive (whereas if $sd^i = 1$ holds at t', it means production of d^i is already completed earlier.)

With this transformation of valid states as well as interactions between them, we first note that, given an interval $i \in \mathcal{I}$ with stochastic demand timing, the optimal decision regarding a regeneration interval $[t_1, t_2]$ is trivial, unless $t_1 \le t_i^* \le t_2 - 1$ holds. Note that there are $O(T^2)$ nontrivial regeneration intervals satisfying $t_1 \le t_i^* \le t_2 - 1$, and for each of these regeneration intervals, we can pre-compute the set of valid vectors $\mathcal{SD}(t_1, t_2)$ as shown above, i.e., by calculating whether it is cheaper to produce d^i at the start or the end of the regeneration interval. With n intervals in total, this would result in at most $O(nT^2)$ computational effort.

Corollary 1. *In the case of time independent production costs and time independent ratio between inventory and backlog costs, the dynamic program has a worst case complexity of $O(nT^2)$.*

4.3. Time independent production costs and convex probability distributions

We next consider the case where the probability distribution for any stochastic demand timing is convex between l_i and u_j. Then, it is straightforward to observe that $EC_i(t)$ is convex, in the same fashion as in Theorem 3 when $\alpha_t = 1\ \forall t$. Therefore, Proposition 2 holds in this case as well, and the worst case complexity of the dynamic program is $O(nT^2)$, as given in Corollary 1.

5. Special cases of stochastic demand timing

In this section, we study two relevant special cases of stochastic demand timing, which enable us to derive very effective dynamic programming algorithms due to the significant reduction of valid states.

5.1. Stochastic demand timing with non-overlapping intervals

First, we consider the case where none of the intervals with stochastic demand timing overlap, i.e., $\forall i, j \in \mathcal{I}$ either $l_i \geq u_j + 1$ or $l_j \geq u_i + 1$ holds. Let the set $\mathcal{I}$ be arranged in an increasing order w.r.t. time, i.e., if $l_i \geq u_j + 1$ for $i, j \in \mathcal{I}$, then $i > j$. For the sake of convenience, we use $G(t, (sd^1, \cdots, sd^n))$ as well as the set of valid vectors at period t, denoted by $\mathcal{SD}(t)$, in line with the general case defined in Section 4.1.

First, let us formalize an important result for this case, which is crucial for the effectiveness of the dynamic program, significantly limiting the state space.

Lemma 2. *For any period t, either $|\mathcal{SD}(t)| = 2$ or $|\mathcal{SD}(t)| = 1$ holds.*

Proof. To observe this result, let us first note that there are three possible cases for any given period t: i) there exists $i \in \mathcal{I}$ such that $t \in [l_i + 1, u_i]$, ii) there exists $i \in \mathcal{I}$ such that $t = l_i$, iii) there does not exist any i satisfying i) or ii). For i), by definition, the only valid vectors are $(sd^j = 1, \forall j \leq i-1; sd^j = 0, \forall j \geq i)$ and $(sd^j = 1, \forall j \leq i; sd^j = 0, \forall j \geq i+1)$, hence $|\mathcal{SD}(t)| = 2$. For ii), by definition, the only valid vector is $(sd^j = 1, \forall j \leq i-1; sd^j = 0, \forall j \geq i)$, hence $|\mathcal{SD}(t)| = 1$. For iii), let $i' = \max\{j \in \mathcal{I} | l_j \leq t-1\}$, i.e., the largest index of the interval starting before t. In this case, there is only one valid vector, which is $(sd^j = 1, \forall j \leq i'; sd^j = 0, \forall j \geq i'+1)$. □

In order to establish the complexity of the dynamic program, we first observe that $n \leq T$ due to the non-overlapping nature of the intervals. Moreover, due to Lemma 2, note that, in the worst case (i.e., when each period is covered by an interval with stochastic demand timing), the state space network will have the same structure with $O(T)$ nodes as the worst case with a single interval discussed in Section 3. Therefore, we conclude this section with the following result.

Corollary 2. *In case of non-overlapping intervals, the dynamic program has a worst case complexity of $O(T \log T)$.*

5.2. Stochastic demand timing with dominant overlapping intervals

In this section, we assume that Property 1 is satisfied for any pair of stochastic demand timings with quantities of d^i and d^j in $\mathcal{I}$, i.e. either d^i dominates d^j or the opposite. Let us also assume that the set $\mathcal{I}$ is arranged in an increasing order w.r.t. to the dominance property, i.e., if d^j is dominated by d^i, then j is ranked after i in $\mathcal{I}$. We next state the key theoretical results for this case.

Theorem 4. *Assume that demand d^i dominates demand d^j, i.e. $\sum_{l=1}^{t} p_l^i \geq \sum_{l=1}^{t} p_l^j$ $\forall t$, and let t^i and t^j be the production periods of d^i and d^j, respectively, in an optimal solution. Then, either i) $t^i \leq t^j$ holds, or ii) both $EC_i(t^i) = EC_i(t^j)$ and $EC_j(t^i) = EC_j(t^j)$ hold.*

In words, the theorem states that in any optimal solution, we will either produce d^i latest in the same period as d^j, or the timing of the production of d^i or d^j is interchangeable between t^i and t^j without any cost implication, i.e., d^i can be produced in t^j, or d^j can be produced in t^i, or both.

Proof. Let $t^i > t^j$ and $EC_i(t^i) \neq EC_i(t^j)$. Since d^i is produced in t^i (rather than in t^j) and $EC_i(t^i) \neq EC_i(t^j)$, $EC_i(t^i) < EC_i(t^j)$ holds, whereas $EC_j(t^i) \geq EC_j(t^j)$ holds since d^j is produced in t^j (rather than in t^i.) By using these relations and equation (5), we have

$$EC_i(t^i) - EC_i(t^j) = -\sum_{l=t^j}^{t^i-1} h_l \sum_{k=t^j+1}^{u_i} p_k^i + \sum_{l=t^j+1}^{t^i-1} b_l \sum_{k=l_i}^{l} p_k^i < 0 \qquad (21)$$

$$EC_j(t^i) - EC_i(t^j) = -\sum_{l=t^j}^{t^i-1} h_l \sum_{k=t^j+1}^{u_j} p_k^j + \sum_{l=t^j+1}^{t^i-1} b_l \sum_{k=l_j}^{l} p_k^j \geq 0 \qquad (22)$$

Then, we have

$$\sum_{l=t^j+1}^{t^i-1} b_l \sum_{k=l_j}^{l} p_k^j \geq \sum_{l=t^j}^{t^i-1} h_l \sum_{k=t^j+1}^{u_j} p_k^j \geq \sum_{l=t^j}^{t^i-1} h_l \sum_{k=t^j+1}^{u_i} p_k^i > \sum_{l=t^j+1}^{t^i-1} b_l \sum_{k=l_i}^{l} p_k^i$$

where the first and third inequalities follow (21) and (22), respectively, and the second inequality follows the dominance property. However, by the dominance property, we also have

$$\sum_{l=t^j+1}^{t^i-1} b_l \sum_{k=l_i}^{l} p_k^i \geq \sum_{l=t^j+1}^{t^i-1} b_l \sum_{k=l_j}^{l} p_k^j$$

which is a contradiction. The same argument follows when $EC_j(t^i) \neq EC_j(t^j)$ holds instead of $EC_i(t^i) \neq EC_i(t^j)$. □

Theorem 5. *Assume that demand d^i dominates demand d^j. If $\exists t$ s.t. $\sum_{l=1}^{t'} p_l^i = \sum_{l=1}^{t'} p_l^j$, $\forall t' \leq t$, and if there is an optimal solution in which d^j is produced before period t, then there is an optimal solution in which both d^i and d^j are produced in the same period.*

Proof. Follows from Theorem 4. □

Finally, we note the following result, which follows from Theorem 5.

Corollary 3. *If, for demands d^i and d^j, $\sum_{l=1}^{t} p_l^i = \sum_{l=1}^{t} p_l^j$ $\forall t$, i.e., they follow exactly the same distribution, then there is an optimal solution in which d^i and d^j are produced in the same period.*

A very useful aspect of Corollary 3 is that stochastic demand timings satisfying these conditions can be merged into a single demand. In the remainder of this subsection, we assume that all such demands are already combined into single demands.

Using $G(t, (sd^1, \cdots, sd^n))$ and the set of valid vectors at period t, denoted by $\mathcal{SD}(t)$, in line with the general case defined in Section 4.1, we formalize the following result for the complexity of the dynamic program.

Lemma 3. *In the worst case,* $\max_{t \in [1,T]} |\mathcal{SD}(t)| \sim O(n)$

This is quite straightforward to observe due to the fact that if $sd^i = 0$ for any time period t, then $sd^j = 0, \forall j \in \mathcal{I}$ s.t. $i < j$. In order to establish the complexity of the dynamic program, we note that in the worst case (i.e., when each period has n states), the state space network will have $O(nT)$ nodes in a structure similar to the case with non-overlapping intervals, albeit with n layers. Hence, we conclude this section with the following result.

Corollary 4. *In case of dominant overlapping intervals, the dynamic program has a worst case complexity of $O(nT \log T)$.*

6. Extension to demand-dependent backlog cost

In this section, we extend our previous results to the more general case where each stochastic demand timing can be seen as an order with a specific backlog cost, i.e. b_t^i now depends both on period t and on stochastic demand quantity d^i. This case naturally stems from the varying importance of satisfying different orders (or, likely, different customers) on time, and provides planners and decision makers a more customized solution.

Next, we discuss the impact of considering b_t^i instead of b_t in order to extend previous results:

- In Section 2, it is straightforward to extend the model by replacing b_l by b_l^i in the definition of $EC_i(t)$, and Proposition 1 remains true.

- Because only a single stochastic demand timing is considered, the analysis and results in Section 3 remain valid in this case as well.
- With respect to the general case of Section 4, it is easy to observe that the general dynamic program proposed in Section 4.1 does not change. The analysis and results of Section 4.2 remain valid when the time independence of the ratio between the unit inventory and backlog costs in each period is redefined as $b_t^i = \alpha_t b^i$ with $\alpha_t > 0$ $\forall t$ and $\forall i$ (or, equivalently, $h_t/b_t^i = h/b^i$, $\forall t$ and $\forall i$). It is also straightforward to see that the discussion on the case with convex probability distributions presented in Section 4.3 remains valid.
- Finally, let us now consider the two special cases studied in Section 5. The analysis conducted in Section 5.1 in the case of non-overlapping intervals remains valid. On the other hand, the notion of dominance between stochastic demand timings, defined in Section 5.2 (Definition 1), is no longer sufficient to derive the results in Section 5.2, which we discuss next.

Definition 3. Let d^i and d^j be two demands with stochastic timing. If $EC_i(t) \geq EC_j(t)$ $\forall t = 1, \ldots, T$, then we say that d^i **strongly dominates** d^j.

In comparison to Definition 1, Definition 3 proposes a stricter definition of dominance, and this can be employed to replace Theorems 4 and 5 with Theorems 6 and 7, and Corollary 3 with Corollary 5. We note that Lemma 3 and Corollary 4 still remain valid.

Theorem 6. *Assume that demand d^i strongly dominates demand d^j, and let t^i and t^j be the production periods of d^i and d^j, respectively, in an optimal solution. Then, either i) $t^i \leq t^j$ holds, or ii) both $EC_i(t^i) = EC_i(t^j)$ and $EC_j(t^i) = EC_j(t^j)$ hold.*

Theorem 7. *Assume that demand d^i strongly dominates demand d^j. If $\exists t$ s.t. $EC_i(t') = EC_j(t')$, $\forall t' \leq t$, and if there is an optimal solution in which d^j is produced before period t, then there is an optimal solution in which both d^i and d^j are produced in the same period.*

Corollary 5. *If, for demands d^i and d^j, $EC_i(t) = EC_j(t)$ $\forall t$, then there is an optimal solution in which d^i and d^j are produced in the same period.*

We omit the proofs of these results, as they can be easily carried out in the same fashion as the proofs of Section 5.2. Finally, we conclude this section with the following important complexity result, which shows that the problem studied in this section is $\mathcal{NP}$-hard.

Theorem 8. *The dynamic lot sizing problem with stochastic demand timing is $\mathcal{NP}$-hard when the backlog cost b_t^i depends both on period t and on the quantity of stochastic demand, d^i.*

Proof. Consider the Uncapacitated Facility Location (UFL) problem, which is $\mathcal{NP}$-hard and can be stated as follows: Given F facilities, each with a fixed opening cost $\alpha_t, t \in F$ and no capacity, and C clients with unit demand and unit service cost of β_{ti} for each client $i \in C$ and each facility $t \in F$, find the subset of facilities to open that will serve all clients with the minimum total cost. Given an instance of the uncapacitated facility location, we will show that this problem can be reduced to an instance of our problem.

First, we create a dummy first period (period 0), and then map facilities of UFL to periods starting from period 1, and map clients of UFL to intervals with stochastic demand timing, hence creating a problem with $F+1$ periods and C intervals, each interval i spanning the whole horizon, i.e., $\ell_i = 0, u_i = F+1$. Then, we can make the following assignments of parameters in our problem:

$$D_t := 0, c_t := 0, \forall t \in [0, F]; f_0 := M; f_t := \alpha_t, \forall t \in [1, F];$$

$$d^i := 1, \forall i \in [1, C]$$

Here, M is a sufficiently big number so that production never takes place in period 0. Note that these assignments result in a dynamic lot sizing problem without z_{lt} variables and instead only with y_t and z_t^i variables. In order to finalize the problem reduction, we make the following assignment:

$$p_t^i := \frac{1}{F+1}, \forall t \in [0, F], \forall i \in [1, C]; \quad EC_i(t) := \beta_{ti}, \forall t \in [1, F], \forall i \in [1, C]$$

Then, specific b_t^i and h_t values can be calculated by solving $F \times C$ linear equations ($\forall t \in [1, F]$, $\forall i \in [1, C]$) of (5) with assigned β_{ti} values and $F \times C + C$ unknowns (b_t^i and h_t, respectively), where the first set of equations for $t = 1$ involve the i-specific parameter b_t^i. It is straightforward to see that a solution of UFL is equivalent to a solution of the single-item dynamic lot sizing problem with stochastic demand timing. □

We make a final remark that the proof relies on the fact that the parameters b_t^i are defined separately for each i. Therefore, the proof is not valid when $b_t^i = b_t$ $\forall i$ as in the original problem, which leaves its complexity open.

7. Conclusions and perspectives

An original and relevant way of modeling stochastic demands in lot sizing problems is proposed and studied in this paper. The uncertainty is not on the demand quantity but rather on the timing at which the demand will occur and should be satisfied. More precisely, the demand quantity of each stochastic demand is known and fully occurs with a given probability in a single period of a given interval and no partial delivery is allowed. In our modeling, stochastic demands are naturally correlated since a stochastic demand occurring in a period will not occur in another period. Dynamic programs are proposed to solve several cases of the single-item dynamic lot sizing problem with stochastic demand timing. The case with a single interval is first solved, followed by the general case, the practical case where the ratio between the inventory and backlog costs is time independent and the case where probability distributions are convex. Finally, the cases with non-overlapping intervals and with dominant overlapping intervals are solved. Note that all the dynamic programs presented in the paper could be extended to the case with backlog costs on deterministic demands and on stochastic demand timings after the last period in their interval. We also study the general case where the variable backlog cost depends both on the period and on the quantity of stochastic demand, and show that the resulting problem is $\mathcal{NP}$-hard.

Many research avenues are worthwhile investigating from this novel stochastic setting in lot sizing. First, although we believe it is $\mathcal{NP}$-hard, the complexity of the general problem with backlog costs that are independent of the quantity of stochastic demand remains an open question to study. Second, the capacitated case with multiple products could be solved using a Lagrangian heuristic, such as the ones proposed in Trigeiro, Thomas, & McClain (1989) and Brahimi, Dauzère-Pérès, & Najid. (2006), by relaxing the capacity constraints and solving the resulting single-item problems with the dynamic programs proposed in this paper. Another interesting extension of our work is to consider the case where $\sum_{t=l_i}^{u_i} p_t^i < 1$, i.e., there is a probability that demand d^i may not occur at all. In this case, the total demand on the planning horizon also becomes uncertain. This implies that some production quantity aimed at satisfying d^i might end up remaining in the inventory and thus could be used to satisfy other demands in the planning horizon. A last related research perspective would be to analyze the case with lost sales, where answering a demand too late would also result in products remaining in the inventory.

Acknowledgements

The work of the first author is supported by the Air Force Office of Scientific Research under award number FA9550-18-1-7003. We are also thankful to two anonymous reviewers for their insightful comments that helped us to significantly improve the content and presentation of the paper.

References

Aggarwal, A., & Park, J. (1993). Improved algorithms for economic lot size problems. *Operations Research, 41*(3), 549–571.
Agra, A., Santos, M. C., Nace, D., & Poss, M. (2016). A dynamic programming approach for a class of robust optimization problems. *SIAM Journal on Optimization, 26*(3), 1799–1823.
Aloulou, M. A., Dolgui, A., & Kovalyov, M. Y. (2014). A bibliography of non-deterministic lot-sizing models. *International Journal of Production Research, 52*(8), 2293–2310.
Attila, Ö. N., Agra, A., Akartunalı, K., & Arulselvan, A. (2021). Robust formulations for economic lot-sizing problem with remanufacturing. *European Journal of Operational Research, 288*(2), 496–510.
Ben-Tal, A., Golany, B., Nemirovski, A., & Vial, J. P. (2005). Retailer-supplier flexible commitments contracts: A robust optimization approach. *Manufacturing & Service Operations Management, 7*(3), 248–271.
Bertsimas, D., & Thiele, A. (2006). A robust optimization approach to inventory theory. *Operations Research, 54*(1), 150–168.
Bienstock, D., & Özbay, N. (2008). Computing robust basestock levels. *Discrete Optimization, 5*(2), 389–414.
Brahimi, N., Absi, N., Dauzère-Pérès, S., & Nordli, A. (2017). Single-item dynamic lot-sizing problems: An updated survey. *European Journal of Operational Research, 263*(3), 838–863.
Brahimi, N., Dauzère-Pérès, S., & Najid, N. (2006). Capacitated multi-item lot-sizing problems with time windows. *Operations Research, 54*(5), 951–967.
Christiansen, M., Fagerholt, K., Flatberg, T., Haugen, Ø., Kloster, O., & Lund, E. H. (2011). Maritime inventory routing with multiple products: A case study from the cement industry. *European Journal of Operational Research, 208*(1), 86–94.
Dauzère-Pérès, S., Brahimi, N., Najid, N., & Nordli, A. (2002). The single-item lot sizing problem with time windows. *Technical Report, 02/4/AUTO*.
Dauzère-Pérès, S., Nordli, A., Olstad, A., Haugen, K., Koester, U., Per Olav, M., … Reistad, A. (2007). Omya hustadmarmor optimizes its supply chain for delivering calcium carbonate slurry to european paper manufacturers. *Interfaces, 37*(1), 39–51.
Doostmohammadi, M., & Akartunalı, K. (2018). Valid inequalities for two-period relaxations of big-bucket lot-sizing problems: Zero setup case. *European Journal of Operational Research, 267*(1), 86–95.
Federgruen, A., & Heching, A. (1999). Combined pricing and inventory control under uncertainty. *Operations Research, 47*(3), 454–475.
Federgruen, A., & Tzur, M. (1993). The dynamic lot-sizing model with backlogging: A simple O*nlogn* algorithm and minimal forecast horizon procedure. *Naval Research Logistics, 40*(4), 459–478.
Golari, M., Fan, N., & Jin, T. (2017). Multistage stochastic optimization for production-inventory planning with intermittent renewable energy. *Production and Operations Management, 26*(3), 409–425.
Gorissen, B. L., Yanikoğlu, İ., & den Hertog, D. (2015). A practical guide to robust optimization. *Omega, 53*, 124–137.
Guan, Y., & Miller, A. J. (2008). Polynomial-time algorithms for stochastic uncapacitated lot-sizing problems. *Operations Research, 56*(5), 1172–1183.
van Hoesel, C. (1991). *Models and algorithms for single item lot-sizing problems*. Erasmus University Rotterdam, Rotterdam, The Netherlands Ph.D. thesis..
Huang, K., & Ahmed, S. (2010). A stochastic programming approach for planning horizons of infinite horizon capacity planning problems. *European Journal of Operational Research, 200*(1), 74–84.
Huang, K., & Küçükyavuz, S. (2008). On stochastic lot-sizing problems with random lead times. *Operations Research Letters, 36*(3), 303–308.
Jiang, R., & Guan, Y. (2011). An $O(N^2)$-time algorithm for the stochastic uncapacitated lot-sizing problem with random lead times. *Operations Research Letters, 39*(1), 74–77.
Lee, C., Çetinkaya, S., & Wagelmans, A. (2001). A dynamic lot-sizing model with demand time windows. *Management Science, 47*(10), 1384–1395.
Levi, R., & Shi, C. (2013). Approximation algorithms for the stochastic lot-sizing problem with order lead times. *Operations Research, 61*(3), 593–602.
Meistering, M., & Stadtler, H. (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.
Tempelmeier, H. (2007). On the stochastic uncapacitated dynamic single-item lot-sizing problem with service level constraints. *European Journal of Operational Research, 181*(1), 184–194.
Tempelmeier, H. (2013). Stochastic lot sizing problems. In *Handbook of stochastic models and analysis of manufacturing system operations* (pp. 313–344). Springer.
Thomas, L. J. (1974). Price and production decisions with random demand. *Operations Research, 22*(3), 513–518.
Trigeiro, W. W., Thomas, L. J., & McClain, J. O. (1989). Capacitated lot sizing with setup times. *Management Science, 35*(3), 353–366.
Wagelmans, A., van Hoesel, S., & Kolen, A. (1992). Economic lot sizing: An $O(n \log n)$ algorithm that runs in linear time in the Wagner-Whitin case. *Operations Research, 40*(1), S145–S156.
Wagner, H., & Whitin, T. (1958). Dynamic version of the economic lot size model. *Management Science, 5*(1), 89–96.
Wei, C., Li, Y., & Cai, X. (2011). Robust optimal policies of production and inventory with uncertain returns and demand. *International Journal of Production Economics, 134*(2), 357–367.
Zhang, Y., Shen, Z. M., & Son, S. (2016). Distributionally robust optimization of two-stage lot-sizing problems. *Production and Operations Management, 25*(12), 2116–2131.