

Computational Complexity of Finding Pareto Efficient Outcomes for Biobjective Lot-Sizing Models

H. Edwin Romeijn,¹ Dolores Romero Morales,² Wilco Van den Heuvel³

¹*Department of Industrial and Operations Engineering, University of Michigan, Ann Arbor, Michigan 48109-2117*

²*Saïd Business School, University of Oxford, Oxford OX1 1HP, United Kingdom*

³*Econometric Institute, Erasmus University Rotterdam, 3000 DR Rotterdam, The Netherlands*

Received 4 June 2013; revised 8 June 2014; accepted 9 June 2014

DOI 10.1002/nav.21590

Published online 4 July 2014 in Wiley Online Library (wileyonlinelibrary.com).

Abstract: In this article, we study a biobjective economic lot-sizing problem with applications, among others, in green logistics. The first objective aims to minimize the total lot-sizing costs including production and inventory holding costs, whereas the second one minimizes the maximum production and inventory block expenditure. We derive (almost) tight complexity results for the Pareto efficient outcome problem under nonspeculative lot-sizing costs. First, we identify nontrivial problem classes for which this problem is polynomially solvable. Second, if we relax any of the parameter assumptions, we show that (except for one case) finding a single Pareto efficient outcome is an \mathcal{NP} -hard task in general. Finally, we shed some light on the task of describing the Pareto frontier. © 2014 Wiley Periodicals, Inc. *Naval Research Logistics* 61: 386–402, 2014

Keywords: lot-sizing; biobjective; expenditure; Pareto efficient outcomes; complexity analysis

1. INTRODUCTION

In this article, we model the choice of a production plan as a biobjective economic lot-sizing problem (BOLS). This is a generalization of the economic lot-sizing (ELS) problem introduced by [27], in which we aim to minimize two objectives. The first objective is the standard objective in the classical ELS problem. It accounts for the total lot-sizing costs across the full planning horizon, of length T , of fixed and variable production, as well as linear inventory holding costs. The second objective accounts for the maximum expenditure among consecutive and disjoint blocks of length ℓ , $1 \leq \ell \leq T$. Here expenditure again refers to fixed and variable production, as well as linear inventory costs; however, possibly evaluated with different cost parameters than in the “total cost” objective.

There are different situations in which this second objective is required in lot-sizing. The second objective may be used to model the processing of a scarce resource for which the maximum available capacity is limited physically or due to regulation. The latter is the case in terms of the carbon footprint of

lot-sizing, where carbon emissions may be released by any of the three processes involved, namely the setup of the machinery, the functioning of the machinery during production, as well as from holding inventory, see [3]. In this first application, BOLS can be used to describe the tradeoff between lot-sizing costs and carbon emissions. The second objective may also be helpful when a resource needs to be balanced across the blocks. In the second application of BOLS, we aim to spread out the lot-sizing costs across the blocks, yielding a cost-balanced plan. Here, both objectives refer to lot-sizing costs, but the first one addresses the full planning horizon, whereas the second addresses each block. In the third application, we are concerned with one of the three processes, per se, say the total production quantity in each block, and we wish to balance the plan with respect to the production quantities.

There are some recent studies in the literature on lot-sizing models with carbon emission constraints. These are hence related to our first application. Ref. [3] is among the first to include carbon emissions in the classical ELS problem, by means of emission caps, taxes on emissions, cap-and-trade emission mechanisms, or carbon offsets. Using a collection of instances, and for each of these four policies, the authors illustrate the impact of the lot-sizing decisions on carbon emissions. Besides the fact that Ref. [3] does not study the

Correspondence to: Dolores Romero Morales (dolores.romero-morales@sbs.ox.ac.uk)

block model, their focus is on managerial insights, in contrast to our contribution, which is of a methodological nature. Ref. [25] provides an empirical study comparing different levels of data aggregation when estimating carbon transportation emissions. The problem is formulated as a lot-sizing model with a carbon emission constraint, and its computational complexity is left as an open question. Refs. [19] and [1] provide theoretical and algorithmic results on lot-sizing models with carbon emission constraints. In [19], a single emission constraint for the whole horizon is considered, where the costs and emissions are modeled by general concave functions of the production values. After showing its \mathcal{NP} -hardness, the authors develop a Lagrangian heuristic and a fully polynomial-time approximation scheme (FPTAS) for the problem. Ref. [1] considers multiple production modes, each of them with its own production cost function and unit carbon production emission parameter. The lot-sizing problem consists of choosing a combination of these modes in each period such that the total production and inventory holding costs are minimized, while a cap is imposed on carbon emissions. Four different ways of imposing this unit cap are proposed: periodic, cumulative, global, and rolling. The authors prove that the periodic model is polynomially solvable, whereas the rest of the models are \mathcal{NP} -hard.

If more than one objective function is optimized, Pareto efficient outcomes (in the value space) are sought. The collection of Pareto efficient outcomes, the Pareto frontier, can be used to describe the tradeoff between the different objectives. In particular, for BOLS, it is of interest to know when the expenditure has a great impact on lot-sizing costs. Pareto efficient outcomes can be found by minimizing one objective function while constraining the others. For multiobjective combinatorial optimization problems, such as BOLS, finding a single Pareto efficient outcome is, in general, an \mathcal{NP} -hard task [8, 9, 22]. The goal of this article is to find classes of instances for which the Pareto efficient outcome problem, in which we minimize the lot-sizing costs and restrict the block expenditure, is polynomially solvable. In particular, we analyze classes for which the costs are nonspeculative (i.e., it is better to postpone production as much as possible in terms of variable costs).

The remainder of the article is organized as follows. In Section 2, we introduce the BOLS and its applications. In Section 3, we formulate the Pareto efficient outcome problem. In Sections 4–6, we will separately investigate the computational complexity of the Pareto efficient outcome problem for whole horizon expenditure ($\ell = T$), period by period expenditure ($\ell = 1$), and block expenditure (general ℓ). In each section, we identify \mathcal{NP} -hard and polynomially solvable classes of instances and discuss the tightness of the complexity results. In Section 7, we shed some light on the task of describing the Pareto frontier. We conclude the article in Section 8 with a summary and topics for future research.

2. THE BOLS MODEL

Consider a planning horizon of length T . For period t , $t = 1, \dots, T$, let f_t be the setup cost, c_t the unit production cost, and h_t the unit inventory holding cost. Similarly, for period t , let \hat{f}_t be the setup expenditure, \hat{c}_t the unit production expenditure, and \hat{h}_t the unit inventory holding expenditure. Let d_t be the demand in period t , with $d_{s,t} = \sum_{j=s}^t d_j$, and M a constant such that $M \geq d_{1,T}$. For ease of notation, an interval of consecutive periods $\{s, \dots, t\}$ is denoted by $[s, t]$.

Let us partition the time horizon into consecutive blocks of ℓ periods. Without loss of generality, we assume that ℓ divides T . If not, we can define an equivalent problem by adding “dummy” periods with demand equal to zero. We are interested in minimizing the costs as well as the maximum expenditure among the blocks of length ℓ . The BOLS model with block length ℓ , hereafter (BOLS^(ℓ)), reads as follows:

$$\begin{aligned} & \text{minimize} \left(\sum_{t=1}^T [f_t y_t + c_t x_t + h_t I_t], \right. \\ & \left. \max_{i=1, \dots, T/\ell} \left\{ \sum_{t=(i-1)\ell+1}^{i\ell} [\hat{f}_t y_t + \hat{c}_t x_t + \hat{h}_t I_t] \right\} \right) \end{aligned}$$

subject to (BOLS^(ℓ))

$$x_t + I_{t-1} = d_t + I_t \quad t = 1, \dots, T \tag{1}$$

$$x_t \leq M y_t \quad t = 1, \dots, T \tag{2}$$

$$I_0 = 0 \tag{3}$$

$$y_t \in \{0, 1\} \quad t = 1, \dots, T$$

$$x_t \geq 0 \quad t = 1, \dots, T$$

$$I_t \geq 0 \quad t = 1, \dots, T,$$

where y_t indicates whether a setup has been placed in period t , x_t denotes the quantity produced in period t , and I_t denotes the inventory level at the end of period t . In the following, we will refer to a production period as a period in which a positive amount of production occurs, that is, $x_t > 0$. The first objective in (BOLS^(ℓ)) models the usual lot-sizing costs, that is, the fixed-charge and variable production, as well as linear inventory holding costs over the whole planning horizon. The second objective function models the maximum block expenditure. Constraints (1) model the balance between production, inventory, and demand in period t . Constraints (2) ensure that the production level is equal to zero if no setup is incurred in period t . Constraint (3) states that the inventory level is equal to zero at the beginning of the planning horizon. The last three constraints define the range on which the variables are defined. When convenient, we will use y , x , and I to refer to the vectors (y_t) , (x_t) , and (I_t) .

The values of the expenditure parameters \hat{f}_t , \hat{c}_t , and \hat{h}_t depend on the area of application of (BOLS^(ℓ)). In the carbon

footprint setting, these would normally be time invariant, that is, $\hat{f}_t = \hat{f}$, $\hat{c}_t = \hat{c}$, and $\hat{h}_t = \hat{h}$ for all t . In the second application, where we aim to balance the lot-sizing costs across the blocks, the expenditure, and the cost parameters coincide, that is, $\hat{f}_t = f_t$, $\hat{c}_t = c_t$, and $\hat{h}_t = h_t$. Note that this application is only of interest if the block length is not the full one, that is, $\ell < T$. Finally, when balancing one of the three processes all the expenditure parameters equal zero, except for the ones associated with the process in question, which are equal to 1. For instance, when balancing production quantities, we would have $\hat{f}_t = 0$, $\hat{c}_t = 1$, and $\hat{h}_t = 0$ for all t .

3. PARETO EFFICIENT OUTCOMES

Given $\hat{b} \in \mathbb{R}_+$, the following problem, hereafter $(\mathcal{P}^{(\ell)}(\hat{b}))$, defines a Pareto efficient outcome for $(\text{BOLS}^{(\ell)})$:

$$z(\hat{b}) = \text{minimize } \sum_{t=1}^T [f_t y_t + c_t x_t + h_t I_t]$$

subject to ($\mathcal{P}^{(\ell)}(\hat{b})$)

$$\begin{aligned} x_t + I_{t-1} &= d_t + I_t \quad t = 1, \dots, T \\ x_t &\leq M y_t \quad t = 1, \dots, T \\ I_0 &= 0 \\ y_t &\in \{0, 1\} \quad t = 1, \dots, T \\ x_t &\geq 0 \quad t = 1, \dots, T \\ I_t &\geq 0 \quad t = 1, \dots, T \end{aligned}$$

$$\sum_{t=(i-1)\ell+1}^{i\ell} [\hat{f}_t y_t + \hat{c}_t x_t + \hat{h}_t I_t] \leq \hat{b} \quad i = 1, \dots, T/\ell. \quad (4)$$

To obtain a Pareto efficient outcome when $\ell = T$, Constraints (4) become the single constraint

$$\sum_{t=1}^T [f_t y_t + c_t x_t + h_t I_t] \leq \hat{b}, \quad (5)$$

resulting in model $(\mathcal{P}^{(T)}(\hat{b}))$. Similarly, when $\ell = 1$, Constraints (4) become

$$\hat{f}_t y_t + \hat{c}_t x_t + \hat{h}_t I_t \leq \hat{b} \quad t = 1, \dots, T, \quad (6)$$

resulting in model $(\mathcal{P}^{(1)}(\hat{b}))$. For ease of reference, we will call models $(\mathcal{P}^{(\ell)}(\hat{b}))$, $(\mathcal{P}^{(T)}(\hat{b}))$, and $(\mathcal{P}^{(1)}(\hat{b}))$ the block model, the whole horizon model and the period model, respectively. Models $(\mathcal{P}^{(T)}(\hat{b}))$ and $(\mathcal{P}^{(1)}(\hat{b}))$ with the addition of backlogging can be found in [3].

If the expenditure constraints are not binding, $(\mathcal{P}^{(\ell)}(\hat{b}))$ reduces to an ELS problem. When the lot-sizing cost function is concave, the ELS problem is solvable in polynomial

time in the length of the planning horizon T [27, 29]. More efficient algorithms for special cases have been developed in [2, 11, 26].

Furthermore, there are polynomial time algorithms for some ELS problems that can be viewed as special cases of $(\mathcal{P}^{(\ell)}(\hat{b}))$. In particular, there exist polynomial time algorithms when we limit the production in each period by time-invariant capacities [12, 23], and thus $(\mathcal{P}^{(1)}(\hat{b}))$ with only time-invariant production expenditures can be solved in polynomial time. Similarly, the ELS problem with upper bounds on the inventory levels is polynomially solvable [16, 21], and so is $(\mathcal{P}^{(1)}(\hat{b}))$ with only inventory expenditure. In addition, under time-invariant setup cost, the parametric study in [24] provides a polynomial algorithm when we limit the number of setups across the planning horizon.

While the ELS problem and some of its generalizations are polynomially solvable, $(\mathcal{P}^{(\ell)}(\hat{b}))$ is \mathcal{NP} -hard in general, since the capacitated lot-sizing problem (CLSP) with time-variant production capacities is a special case of it [13, 4].

In the following sections, we will investigate which classes of instances of $(\mathcal{P}^{(\ell)}(\hat{b}))$ are \mathcal{NP} -hard or polynomially solvable, where we focus on classes with nonspeculative costs, that is, $c_{t-1} + h_{t-1} \geq c_t$ for $t = 2, \dots, T$. We start with the whole horizon model in Section 4 and then proceed to the period model and block model in Sections 5 and 6, respectively. Each section starts with the identification of classes of instances that are \mathcal{NP} -hard. Subsequently, we develop decomposition algorithms for classes of instances that can be solved in polynomial time. Finally, each section ends with a discussion on the tightness of the complexity results. In the remainder of the article, and when calculating optimal lot-sizing costs, any object in the decomposition with infinite cost corresponds to an infeasible (partial) solution and provides a certificate of infeasibility if required.

We end this section with two remarks on problem $(\mathcal{P}^{(\ell)}(\hat{b}))$. The first remark is on its feasibility. It is clear that $(\mathcal{P}^{(\ell)}(\hat{b}))$ is not feasible if \hat{b} is sufficiently small, say for $\hat{b} < \hat{b}^{\min}$. In case $\ell = T$, the value \hat{b}^{\min} can be found in polynomial time by solving a single unconstrained ELS problem with the total expenditures as the objective function. However, in case $\ell < T$, a similar approach (i.e., solving an unconstrained ELS problem for each block) does not work. The reason is that, in general, the zero inventory order (ZIO) property does not hold at the end of each block, and hence the (minimum) expenditures of one block depend on the (minimum) expenditures of other blocks. However, we can find the value \hat{b}^{\min} to any desired precision by performing a binary search (so solving $(\mathcal{P}^{(\ell)}(\hat{b}))$ repeatedly for different values of \hat{b}), where ‘‘feasibility’’ is the criterion in the search.

The second remark is on the uniqueness of the solution to $(\mathcal{P}^{(\ell)}(\hat{b}))$ with objective value $z(\hat{b})$. In case, the solution is unique, we have found a strongly Pareto efficient outcome. However, in case of multiple solutions, the solution is a

weakly Pareto efficient outcome in general. To get the corresponding strongly Pareto efficient outcome, we need to find the value \hat{b}_L such that $z(\hat{b}_L) = z(\hat{b})$ and $z(\hat{b}') > z(\hat{b})$ for $\hat{b}' < \hat{b}_L$. Based on this observation, the value \hat{b}_L and hence the strongly Pareto efficient outcome can again be found by binary search to any precision. Clearly, a necessary condition to be able to perform both binary searches in a computationally efficient way is the existence of an efficient algorithm to solve $(\mathcal{P}^{(\ell)}(\hat{b}))$, which is the topic of the next sections. How to find the exact values \hat{b}^{\min} and \hat{b}_L in polynomial time is left for future research.

4. WHOLE HORIZON MODEL

4.1. \mathcal{NP} -Hard Instances

Proposition 4.1 shows that the whole horizon model is \mathcal{NP} -hard, even if the setup and the production cost, as well as the inventory holding expenditures, are time-invariant.

PROPOSITION 4.1: Problem $(\mathcal{P}^{(T)}(\hat{b}))$ is \mathcal{NP} -hard under time-invariant setup cost, no production cost, and

- production expenditures only, or
- setup expenditures only, or
- time-invariant holding expenditures only.

PROOF: To show the result for cases (i) and (ii), we use a reduction from the well-known partition problem, which is \mathcal{NP} -complete (see problem [SP12] in [14, p. 223]). We will prove case (i) and only provide the reduced instances for case (ii) since the proof is similar. For case (iii), we use a reduction from the well-known KNAPSACK problem, which is \mathcal{NP} -complete (see problem [MP9] in [14, p. 247]).

CASE (i): As mentioned above, we use a reduction from the PARTITION problem.

Problem PARTITION: Given a set of positive integers $\{a_1, a_2, \dots, a_n\}$, does there exist a set $S \subset N = \{1, \dots, n\}$ with the complement set $S^c = N \setminus S$ such that

$$\sum_{i \in S} a_i = \sum_{i \in S^c} a_i = \frac{1}{2} \sum_{i \in N} a_i = A?$$

We construct an instance for problem $(\mathcal{P}^{(T)}(\hat{b}))$ in polynomial time as follows. Let $T = 2n$, $\hat{b} = A$, and the other parameters according to the following table for $i = 1, \dots, n$:

t	d_t	f_t	c_t	h_t	\hat{f}_t	\hat{c}_t	\hat{h}_t
$2i - 1$	0	A	0	a_i	0	0	0
$2i$	1	A	0	∞	0	a_i	0

Note that the instance can be viewed in terms of pairs of consecutive odd and even periods $(2i - 1, 2i)$. We will show that the answer to partition is yes if and only if there exists a solution with lot-sizing cost at most $(n + 1)A$.

Suppose the answer to partition is yes, that is, there exists a set S such that $\sum_{i \in S} a_i = A$. We construct a solution to the lot-sizing instance as follows: $y_t = x_t = 1$ when $t = 2i - 1$ with $i \in S$, or when $t = 2i$ with $i \in S^c$. In other words, every demand is satisfied from either the current or the previous period. Since there are n setups, the total setup costs equal nA . Furthermore, the total holding costs equal $\sum_{i \in S} a_i = A$, while the total expenditure equals $\sum_{i \in S^c} a_i = A$. Therefore, the solution is feasible and has costs $(n + 1)A$.

Now suppose that there exists a feasible solution (x, y) with costs at most $(n + 1)A$. Because of the holding cost structure, demand in period $2i$ is satisfied from period $2i - 1$ or $2i$. Moreover, demand will be satisfied from a single period since otherwise the costs will exceed $(n + 1)A$ or the solution is infeasible. This means that there are n setups incurring a cost of nA . Let S (resp. T) be the set of indices for which demand is satisfied by the odd (resp. even) periods, that is, $S = \{i : y_{2i-1} = 1\}$ (resp. $T = \{i : y_{2i} = 1\}$). Clearly, (S, T) is a partition of N and hence $T = S^c$. Therefore, the total holding costs equal $\sum_{i \in S} a_i$, while the total expenditure equals $\sum_{i \in S^c} a_i$. Since the sum of the two is $2A$, we must have $\sum_{i \in S} a_i = \sum_{i \in S^c} a_i = A$ in a feasible solution with costs at most $(n + 1)A$. Hence, the sets S and S^c give the desired partition. (Note that the proof also holds if we let $d_t = a_i$ for $t = 2i$ and replace the values of h_t for the odd periods with 1 as well as the values \hat{c}_t in the even periods.)

CASE (ii): Here, again, we use a reduction from the partition problem. The reduced instance has $T = 2n$, $\hat{b} = A$, and the other parameters according to the following table for $i = 1, \dots, n$:

T	d_t	f_t	c_t	h_t	\hat{f}_t	\hat{c}_t	\hat{h}_t
$2i - 1$	0	A	0	a_i	0	0	0
$2i$	1	A	0	∞	a_i	0	0

(Note that the result still holds if we let $d_t = a_i$ for $t = 2i$ and replace the finite values of h_t with 1.)

CASE (iii): For this case, we use a reduction from the KNAPSACK problem.

Problem KNAPSACK: Given positive integers $\kappa, \kappa', a_1, a_2, \dots, a_n, a'_1, a'_2, \dots, a'_n$, does there exist a vector $z \in \{0, 1\}^n$ such that

$$\sum_{i=1}^n a_i z_i \geq \kappa \quad \text{and} \quad \sum_{i=1}^n a'_i z_i \leq \kappa'?$$

We construct an instance for problem $(\mathcal{P}^{(T)}(\hat{b}))$ in polynomial time as follows. Let $T = 2n$, $\hat{b} = \kappa'$, and the other parameters according to the following table for $i = 1, \dots, n$:

t	d_t	f_t	c_t	h_t	\hat{f}_t	\hat{c}_t	\hat{h}_t
$2i - 1$	1	A	0	$\frac{A-a_i}{a_i}$	0	0	1
$2i$	a'_i	A	0	∞	0	0	1

where A is a sufficiently large number. As in cases (i) and (ii), the instance can be viewed in terms of pairs of consecutive odd and even periods $(2i - 1, 2i)$, where d_{2i-1} is always produced in period $2i - 1$. We will show that the answer to KNAPSACK is yes if and only if there exists a solution with lot-sizing cost at most $2nA - \kappa$.

Suppose the answer to KNAPSACK is yes, that is, there exists a vector $z \in \{0, 1\}^n$ such that $\sum_{i=1}^n a_i z_i \geq \kappa$, $\sum_{i=1}^n a'_i z_i \leq \kappa'$. We construct a solution to the lot-sizing instance as follows. For $i = 1, \dots, n$ and $t = 2i - 1$, we have $y_t = 1$. For $i = 1, \dots, n$ and $t = 2i$, $y_t = 0$ and $I_{t-1} = d_t$ if $z_i = 1$, and $y_t = 1$ and $I_{t-1} = 0$ otherwise. In other words, we always produce in period $2i - 1$, but in addition we will produce in period $2i$ if $z_i = 0$. The total setup costs equal $An + \sum_{i=1}^n A(1 - z_i) = 2nA - \sum_{i=1}^n Az_i$. Furthermore, the total holding costs equal $\sum_{i=1}^n \frac{A-a_i}{a_i} a'_i z_i = \sum_{i=1}^n (A - a_i) z_i$, while the total expenditure equals $\sum_{i=1}^n a'_i z_i \leq \kappa'$. Therefore, the solution is feasible and has costs $2nA - \sum_{i=1}^n Az_i + \sum_{i=1}^n (A - a_i) z_i = 2nA - \sum_{i=1}^n a_i z_i \leq 2nA - \kappa$.

Now suppose that there exists a feasible solution (y, I) with costs at most $2nA - \kappa$. Note that we can assume that demand in period $2i$ will be satisfied either from period $2i - 1$ or $2i$, but not both. Otherwise, we can decrease I_{2i-1} to zero, reducing thus the inventory holding expenditure as well as the inventory holding costs. For each $i = 1, \dots, n$, we can now define $z_i = 1$ if $y_{2i} = 0$ and $z_i = 0$ otherwise. Due to the structure of the inventory holding cost, $I_{2i} = 0$ for all i , while $I_{2i-1} = d_{2i}$ if $z_i = 0$ and 0 otherwise. Therefore, the expenditure of the lot-sizing solution is equal to $\sum_{i=1}^n d_{2i} z_i = \sum_{i=1}^n a'_i z_i$ and therefore at most κ' . With similar manipulations as above, we can show that the lot-sizing costs can be written as

$$2nA - \sum_{i=1}^n a_i z_i$$

and therefore $\sum_{i=1}^n a_i z_i \geq \kappa$. □

Note that, to find a Pareto efficient solution of $(\text{BOLS}^{(T)})$, one could optimize over the second objective and constrain the first objective as well (instead of optimizing over the first and constraining the second objective). Moreover, if we subsequently interchange the role of cost and expenditure, the problem reduces to $(\mathcal{P}^{(T)}(\hat{b}))$ again. That means that the

complexity results for problem $(\mathcal{P}^{(T)}(\hat{b}))$ are symmetric with respect to the parameter assumptions on cost and expenditure. Therefore, the following corollary is immediate from Proposition 4.1.

COROLLARY 4.2: Problem $(\mathcal{P}^{(T)}(\hat{b}))$ is \mathcal{NP} -hard under time-invariant setup expenditures, no production expenditure and

1. production cost only, or
2. setup cost only, or
3. time-invariant holding cost only.

4.2. Polynomially Solvable Instances

In this section, we show that $(\mathcal{P}^{(T)}(\hat{b}))$ can be solved in $\mathcal{O}(T^2)$ time if we have time-invariant f_t , c_t , \hat{f}_t , and \hat{c}_t , and $\hat{h}_t = \alpha h_t$ for some $\alpha \geq 0$. Notice that under these assumptions, the lot-sizing costs and the expenditure are such that there are no speculative motives to hold inventory. Two observations are given before we present the procedure to solve $(\mathcal{P}^{(T)}(\hat{b}))$. First, for a production plan with n production periods, and because both the setup and the unit production expenditure are time-invariant, constraint (5) can be written as

$$\alpha \sum_{t=1}^T h_t I_t \leq \hat{b} - \hat{f}n - \hat{c} \sum_{t=1}^T d_t. \tag{7}$$

Second, and similarly as above, the objective function of $(\mathcal{P}^{(T)}(\hat{b}))$ boils down to

$$\sum_{t=1}^T (f_t y_t + c_t x_t + h_t I_t) = fn + c \sum_{t=1}^T d_t + \sum_{t=1}^T h_t I_t.$$

Thus, when the number of production periods is fixed, minimizing the total inventory costs is equivalent to minimizing the total costs, as well as minimizing the total expenditure. Note that based on the above observations, it is also clear that the ZIO property holds in an optimal solution.

To solve $(\mathcal{P}^{(T)}(\hat{b}))$, it is now sufficient to solve the ELS problem with n production periods for $n = 1, \dots, T$. For a given n , the solution is kept if the inventory levels of the optimal solution satisfy (7). After evaluating all possible values of n , we will have at most T solutions, from which we choose the one having the lowest lot-sizing costs. Solving the ELS problem with n production periods can be done by a straightforward $\mathcal{O}(T^3)$ dynamic programming (DP) algorithm similar to [15], using the variables $F_n(t)$, $n = 1, \dots, T$, $t = n, \dots, T$, where $F_n(t)$ represents the optimal cost of the ELS problem consisting of periods $1, \dots, t$ with n ($n \leq t$) production periods. However, Ref. [24] shows that all values $F_n(t)$ can be found in $\mathcal{O}(T^2)$ time in case of nonspeculative

Table 1. Complexity overview for the whole horizon model.

Costs		Expenditures			Complexity
f_t	c_t and h_t	\hat{f}_t	\hat{c}_t	\hat{h}_t	
c	Nonspeculative	c	c	αh_t	Polynomially solvable
c	Nonspeculative	v	0	0	\mathcal{NP} -hard
c	Nonspeculative	0	v	0	\mathcal{NP} -hard
c	Nonspeculative	0	0	c	\mathcal{NP} -hard

motives to hold inventory, while [10, Section 3.1] obtain the same result in case of time-invariant setup cost. Both results imply that $(\mathcal{P}^{(T)}(\hat{b}))$ under the given parameter assumptions can be solved in $\mathcal{O}(T^2)$ time.

4.3. Discussion of the Complexity Results

In Sections 4.1 and 4.2, we have derived \mathcal{NP} -hardness results and developed polynomial time algorithms for certain classes of instances. A natural question is whether the complexity results are tight, that is, whether each class of instances can be classified as \mathcal{NP} -hard or polynomially solvable. In this section, we show that the answer to this question is positive.

Table 1 gives an overview of the complexity results for the whole horizon model. We use the labels “0,” “c,” “v,” which stand for zero, time-invariant (constant), or time-variant (variable), to specify the assumptions on the parameters. Furthermore, instances with constant (or zero) production cost and time-variant holding cost are labeled as nonspeculative, since it is well-known that any such instance is equivalent to an instance with nonspeculative motives (see e.g., [28]). Suppose that we relax any of the expenditure assumptions for the polynomially solvable instances, that is, (i) time-invariant setup expenditures become time-variant setup expenditures, (ii) time-invariant production expenditures become time-variant expenditures, or (iii) zero holding expenditures ($\alpha = 0$) become time-invariant holding expenditures. Then in each of the three cases it follows that the problem becomes \mathcal{NP} -hard. This means that the results for the whole horizon model are tight.

5. PERIOD MODEL

5.1. \mathcal{NP} -Hard Instances

The \mathcal{NP} -hardness of $(\mathcal{P}^{(1)}(\hat{b}))$ follows from that of the CLSP [13, 4], which is formalized in the following proposition.

PROPOSITION 5.1: Problem $(\mathcal{P}^{(1)}(\hat{b}))$ is \mathcal{NP} -hard under time-invariant setup cost, nonincreasing production cost, zero inventory holding cost and

- production expenditures only, or
- setup expenditures and time-invariant production expenditures only.

PROOF: We will prove case (i) and only provide the reduced instance for case (ii) since the proof is similar.

CASE (i): By rewriting the expenditure constraint (3) as

$$x_t \leq \tilde{b}_t = \frac{\hat{b}}{\hat{c}_t},$$

it is straightforward to see that the CLSP with time-invariant setup cost, nonincreasing production cost, zero inventory holding cost, and time-variant production capacities is a particular case of problem $(\mathcal{P}^{(1)}(\hat{b}))$. Ref. [4] showed that this problem is \mathcal{NP} -hard using a reduction from the subset sum problem (see problem [SP13] in [14, p. 223]). We refer to [4] for the details of the proof.

CASE (ii): The reduced instance in this case has $d_t = 0$, for $t < T$, and $d_T = \kappa$, $\hat{b} = 1 + \max_{t=1, \dots, T} \{a_t\}$, and the other parameters according to the following table for $t = 1, \dots, T$:

f_t	c_t	h_t	\hat{f}_t	\hat{c}_t	\hat{h}_t
1	$\frac{a_t-1}{a_t}$	0	$\hat{b} - a_t$	1	0

Using these parameters, the remainder of the proof parallels the proof in [4] and hence further details are omitted. □

5.2. Polynomially Solvable Instances

In this section, we show that $(\mathcal{P}^{(1)}(\hat{b}))$ can be solved in $\mathcal{O}(T^2)$ time if the lot-sizing costs are such that there are no speculative motives to hold inventory and $f_t \geq f_{t+1}$, whereas for the expenditures we assume that \hat{f}_t, \hat{c}_t , and \hat{h}_t are time-invariant. To solve the problem, we use the concepts of subplan and regeneration period, which are widely used in the literature on lot-sizing problems. We define a subplan as an interval $[s, t]$ of periods such that the starting and ending periods are regeneration periods (i.e., $I_{s-1} = I_t = 0$), whereas the rest of the inventory levels are strictly positive (i.e., $I > 0$ for $r \in [s, t - 1]$). We will refer to a production period as one with positive production. We will say that a period is tight if the corresponding constraint in (3) is binding. We will say that a period is extreme if it is either tight or it is not a production period.

Since \hat{h}_t are time-invariant, when considering constraint (3), the only relevant periods are those in which production occurs. Suppose that period t has positive inventory and no production, that is, $y_t = x_t = 0$. Let t' be the last production

period such that $t' < t$. Since $I_0 = 0$, this period should exist. Since there is no production in periods $t' + 1, \dots, t$, we have that $I_{t'} \geq I_t$. Moreover, $y_{t'} = 1$ and $x_{t'} > 0$ by definition. Thus,

$$\hat{f}y_t + \hat{c}x_t + \hat{h}I_t = \hat{h}I_t \leq \hat{h}I_{t'} \leq \hat{f}y_{t'} + \hat{c}x_{t'} + \hat{h}I_{t'} \leq \hat{b}.$$

In other words, if Constraint (3) holds for each production period, then it will also hold for each nonproduction period. Moreover, if either \hat{f} or \hat{c} is strictly positive, then period t will not be tight.

To solve $(\mathcal{P}^{(1)}(\hat{b}))$ in polynomial time, we will use a decomposition into subplans. In total, there are $\mathcal{O}(T^2)$ subplans. Below, we will show that the optimal costs of all subplans can be found in $\mathcal{O}(T^2)$ time, and thus $(\mathcal{P}^{(1)}(\hat{b}))$ can be solved in $\mathcal{O}(T^2)$ time too. The approach that we follow is similar to [7]. The main difference is that determining the exact production quantities in our approach poses an additional challenge.

The following proposition shows that, for each subplan, all periods will be extreme except for the first one.

PROPOSITION 5.2: There exists an optimal solution for which, except for the first one, all production periods in a subplan are tight.

PROOF: Notice that it is sufficient to show that if t' and t are two consecutive production periods within the same subplan and $t' < t$, then t is tight. This result follows from the nonspeculative motives of the lot-sizing costs to hold inventory.

Suppose that t is a nontight production period and recall that t' is the last production period before t . We can reduce the production in period t' as well as the inventory levels in periods $t', t' + 1, \dots, t - 1$ by ε , at the same time that we increase the production in period t by ε .

Constraint (6) is still satisfied in period $t', t' + 1, \dots, t - 1$, since we have reduced the corresponding left hand side. Constraint (3) for period t is still satisfied if $\hat{f} + \hat{c}(x_t + \varepsilon) + \hat{h}I_t \leq \hat{b}$, which means that $\varepsilon \leq \varepsilon_t = (\hat{b} - (\hat{f} + \hat{c}x_t + \hat{h}I_t))/\hat{c}$. Furthermore, the production in period t' cannot be negative, which means that $\varepsilon \leq x_{t'}$. Finally, we need to impose that $\varepsilon \leq I_{t-1}$ to make sure that the new inventory levels in $t', t' + 1, \dots, t - 1$ are all nonnegative. Thus, for an appropriate choice of ε , namely $\varepsilon = \min\{x_{t'}, I_{t-1}, \varepsilon_t\}$, this solution is still feasible. Observe that because of the nonspeculative motives assumption of the lot-sizing costs, this solution is also optimal.

If $\varepsilon = x_{t'}$, then the result holds for t' and t since t' is not a production period anymore. If $\varepsilon = I_{t-1}$, then the new inventory level at the end of period $t - 1$ is equal to zero and the subplan decomposes into two new subplans. Finally, if $\varepsilon = \varepsilon_t$, then period t becomes tight, and the desired result holds again for periods t' and t . \square

Below we show that production can only occur if the incoming inventory is not enough to satisfy the current demand.

PROPOSITION 5.3: There exists an optimal solution satisfying $I_{t-1} < d_t$ where t is a production period and $t + 1$ a nonproduction period.

PROOF: If t is a production period and $t + 1$ a nonproduction period with $I_{t-1} \geq d_t$ in an optimal solution, then we can move the production quantity from period t to $t + 1$. This solution is still feasible since the inventory levels will not increase and the unit production expenditures are time-invariant. Because of the assumptions on the cost structure, in each period, the costs will not increase and so the solution is also optimal. Repeating the above procedure leads to a solution satisfying the property. \square

The next proposition will be helpful to determine the tight production periods.

PROPOSITION 5.4: Consider a subplan $[u, v]$ and a period $t, u < t \leq v$, with outgoing inventory I_t and satisfying the properties:

- $\bar{x}_t = (\hat{b} - \hat{f} - \hat{h}I_t)/\hat{c} > 0$,
- $\bar{I}_{t-1} = I_t - \bar{x}_t + d_t > 0$.

Then period t is a tight production period in the subplan with production quantity \bar{x}_t , and incoming inventory equal to \bar{I}_{t-1} .

PROOF: We will prove the result by contradiction. Note that \bar{x}_t is the maximum possible production in period t , and therefore \bar{I}_{t-1} is the minimum inventory level in period $t - 1$. Since $\bar{I}_{t-1} > 0$, if t is a production period then it will not be the first one of the subplan, and by Proposition 5.2, t must be tight. Moreover, the production must be equal to \bar{x}_t and the incoming inventory should be \bar{I}_{t-1} .

Thus, we will assume that t is not a production period and show that this is not possible. Let period $s, u \leq s < t \leq v$, be the last production period before t . This means that $x_j = 0$ for $j = s + 1, \dots, t$ and $I_s = I_t + d_{s+1,t}$. Since s should be feasible in terms of expenditure we have

$$x_s \leq (\hat{b} - \hat{f} - \hat{h}(I_t + d_{s+1,t}))/\hat{c} = \bar{x}_t - (\hat{h}d_{s+1,t})/\hat{c}$$

and therefore

$$\begin{aligned} I_{s-1} &= I_s - x_s + d_s \\ &\geq (I_t + d_{s+1,t}) - (\bar{x}_t - (\hat{h}d_{s+1,t})/\hat{c}) + d_s \\ &= \bar{I}_{t-1} + (\hat{h}d_{s+1,t})/\hat{c} + d_{s,t-1} > d_s. \end{aligned}$$

Now we either have a contradiction with $I_{u-1} = 0$ (in case $s = u$) or a contradiction with Proposition 5.3 (in case $s > u$). \square

We can now use Proposition 5.4 to construct an optimal solution to any nondegenerate subplan $[u, v]$, that is, a subplan that does not decompose into multiple subplans, in a backward way. Assume that we arrive at some period $t > u$, for which I_t is known (note that $I_v = 0$ in the initialization of the procedure). We want to determine x_t and I_{t-1} such that the properties of Proposition 5.4 are satisfied. We consider the following cases:

- $\bar{x}_t \leq 0$: The subplan is infeasible, since Constraint (3) is violated for period t or some period before t . To see this, first notice that since $\bar{x}_t \leq 0$ and $\hat{b} > \hat{f}$, we must have $I_t > 0$. Moreover, there should be at least a production period in the interval $[u, t]$.
- By definition, \bar{x}_t is the maximum production quantity in period t without violating the expenditure constraint, and therefore since $\bar{x}_t \leq 0$, we cannot produce in period t .
- The same holds for any period in $[u, t - 1]$. Indeed, it follows from the proof of Proposition 5.4 that any feasible production quantity in period s ($s < u$) is at most equal to \bar{x}_s . In other words, any period with a positive production amount before period t will violate the expenditure constraint too.
- $\bar{x}_t > 0$ and $\bar{I}_{t-1} \leq 0$: In this case, period t cannot be a tight production period, since either production would be too much (in case $\bar{I}_{t-1} < 0$) or the subplan would be degenerate (in case $\bar{I}_{t-1} = 0$). Therefore, we set $x_t = 0$ and $I_{t-1} = I_t + d_t$.
- $\bar{x}_t > 0$ and $\bar{I}_{t-1} > 0$: By Proposition 5.4, period t is tight. Hence, we set $x_t = \bar{x}_t$ and $I_{t-1} = \bar{I}_{t-1}$.

This procedure is applied for periods $t = v, \dots, u + 1$. If we arrive at period u and $0 < d_u + I_{u+1} \leq \bar{x}_u$, then subplan $[u, v]$ is feasible and nondegenerate with a production quantity equal to $x_u = d_u + I_{u+1}$.

For given periods u and v , the cost of subplan $[u, v]$ can be determined in linear time. Hence, a straightforward implementation would lead to an $\mathcal{O}(T^3)$ time algorithm. However, note that when determining subplan $[1, v]$, we also find subplans $[u, v]$ for $u = 1, \dots, v$. This means that the optimal costs of all subplans can be found in $\mathcal{O}(T^2)$ time.

5.3. Discussion of the Complexity Results

We have summarized the complexity results of the period model in Table 2. Since a lot-sizing problem with nonincreasing production and zero holding cost is by definition equivalent to a lot-sizing problem with nonspeculative motives,

Table 2. Complexity overview in case of period expenditures.

Costs		Expenditures			Complexity
f_t	c_t and h_t	\hat{f}_t	\hat{c}_t	\hat{h}_t	
Nonincreasing	Nonspeculative	c	c	c	Polynomially solvable
c	Nonspeculative	v	c	0	\mathcal{NP} -hard
c	Nonspeculative	0	v	0	\mathcal{NP} -hard
c	Nonspeculative	0	0	v	Open

we have labeled the variable cost of the NP-hard cases from Section 5.1 as nonspeculative. The first three lines of the table show that the complexity results are not tight. That is, we have not been able to identify the complexity status of the problem with time-variant holding expenditures only, which is neither a special case of the polynomially solvable problem nor a generalization of the \mathcal{NP} -hard problems. Hence, we have classified the complexity status of this problem as open.

6. BLOCK MODEL

6.1. \mathcal{NP} -Hard Instances

To identify classes of \mathcal{NP} -hard instances for the block model, we should realize that it is both a generalization of the whole horizon and the period models. Therefore, any class of instances that is \mathcal{NP} -hard for any of these two is also \mathcal{NP} -hard for the block model. In Section 6.5, we will discuss these instances in more detail.

In the next sections, we will show that $(\mathcal{P}^{(\ell)}(\hat{b}))$ can be solved in polynomial time if the lot-sizing costs are such that there are nonspeculative motives, whereas for the expenditures we assume that \hat{f}_t and \hat{c}_t are time-invariant and $\hat{h}_t = 0$. We will first prove that the problem can be solved in $\mathcal{O}(T^2\ell)$ time if we face setup expenditures only (Section 6.2). Then, we will move on to the more general case showing an $\mathcal{O}(T^7/\ell)$ time algorithm (Section 6.3), and improve it to $\mathcal{O}(T^5)$ if we face production expenditures only (Section 6.4). To enhance readability, we focus on the main ideas and present the proofs and other details in the appendix.

6.2. Time-Invariant Setup Expenditures

In this section, we show that $(\mathcal{P}^{(\ell)}(\hat{b}))$ can be solved in $\mathcal{O}(T^2\ell)$ time if \hat{f}_t are time-invariant and $\hat{c}_t = \hat{h}_t = 0$, using a variant of the DP algorithm to solve the ELS problem. First note that the ZIO property holds, as stated in the following proposition.

PROPOSITION 6.1: Consider problem $(\mathcal{P}^{(\ell)}(\hat{b}))$ under the additional condition of nonspeculative motives for the

lot-sizing costs and zero production expenditures. Then, without loss of optimality, the ZIO property holds.

Proposition 6.1, together with the expenditure assumptions, implies that an optimal solution can be found under the ZIO solutions with an upper bound on the number of production periods in a block. Let $\bar{b} = \lfloor \hat{b}/\hat{f} \rfloor$ be the maximum number of production periods in any block. Note that we can focus on $\bar{b} < \ell$, and therefore $\bar{b} \in \mathcal{O}(\ell)$, since otherwise the expenditure constraints will be redundant.

Our approach is based on finding a shortest path in a network with nodes (u, n) , $1 \leq u \leq T$, $1 \leq n \leq \bar{b}$, representing that the n^{th} setup of the corresponding block occurs in period u . There are two types of arcs from periods u to v , for which demand $d_{u,v-1}$ is satisfied by production in period u . For each u and v in the same block and $n < \bar{b}$, we have an arc from (u, n) to $(v, n + 1)$, showing that the number of setups in the block increases by one. Similarly, if u and v are in different blocks, we have an arc from (u, n) to $(v, 1)$, since v is the first setup in its associated block. The number of arcs of the first type is $\mathcal{O}(T\ell\bar{b})$ (each block has $\mathcal{O}(\ell\bar{b})$ arcs and there are $\mathcal{O}(T/\ell)$ blocks), while the number of arcs of the second type is $\mathcal{O}(T\bar{b})$. Because of the interpretation of the arcs, the costs of arcs $((u, n), (v, n + 1))$ and $((u, n), (v, 1))$ can be computed in constant time from arcs $((u + 1, n), (v, n + 1))$ and $((u + 1, n), (v, 1))$, respectively. Hence, the problem can be solved in $\mathcal{O}(T^2\bar{b}) \subseteq \mathcal{O}(T^2\ell)$.

6.3. Time-Invariant Setup and Production Expenditures

In this section, we show that $(\mathcal{P}^{(\ell)}(\hat{b}))$ can be solved in $\mathcal{O}(T^7/\ell)$ time if the lot-sizing costs are such that there are no speculative motives to hold inventory, \hat{f}_t and \hat{c}_t are time-invariant and $\hat{h}_t = 0$. Defining $\tilde{b} = \frac{\hat{b}}{\hat{c}}$ and $\tilde{f} = \frac{\hat{f}}{\hat{c}}$, Constraints (4) can be written as

$$\tilde{f} \sum_{\tau=(i-1)\ell+1}^{i\ell} y_{\tau} + \sum_{\tau=(i-1)\ell+1}^{i\ell} x_{\tau} \leq \tilde{b} \quad i = 1, \dots, T/\ell.$$

If the number of setups $\gamma_i \in \{1, \dots, \ell\}$ in block i is known, the constraints reduce to

$$\sum_{\tau=(i-1)\ell+1}^{i\ell} x_{\tau} \leq \tilde{b} - \gamma_i \tilde{f} \quad i = 1, \dots, T/\ell. \tag{8}$$

To develop a DP approach for $(\mathcal{P}^{(\ell)}(\hat{b}))$, we use the following definition.

DEFINITION 6.2: Two consecutive subplans are called connected if they have production in the same block.

$(\mathcal{P}^{(\ell)}(\hat{b}))$ can be decomposed into a collection of maximal sequences of connected subplans, where a sequence is maximal if no other subplan is connected to it. We naturally have




Figure 1. Illustrating the two-tier decomposition for $T = 16$ and $\ell = 4$.

a two-tier decomposition for $(\mathcal{P}^{(\ell)}(\hat{b}))$, where the planning horizon is split into maximal sequences and each sequence is subsequently split into a collection of subplans. This decomposition is illustrated in Fig. 1 for $T = 16$ and $\ell = 4$, where the time horizon is split into two maximal sequences, namely $[1, 5]$ and $[6, 16]$. Some of our results will be illustrated for sequence $[6, 16]$ and therefore we present there a possible plan. The vertical arcs represent production periods, whereas the horizontal arcs represent positive inventory periods. Note that $t = 5$ will not be a production period, otherwise $[6, 16]$ is not maximal.

We now analyze the running time of our two-tier decomposition approach. Clearly, there are $\mathcal{O}(T^2)$ maximal sequences. Given the optimal costs of these sequences, the optimal collection of them can be found by a straightforward DP algorithm in $\mathcal{O}(T^2)$ time. We will show that finding the optimal cost of a maximal sequence can be done by solving a shortest path problem on a network with $\mathcal{O}(T^4/\ell)$ arcs (Section 6.3.2). Given the $\mathcal{O}(T^2)$ maximal sequences, this would lead to a total number of $\mathcal{O}(T^6/\ell)$ arcs. However, since networks corresponding to different maximal sequences share identical arcs, the total number of different arcs is $\mathcal{O}(T^5/\ell)$ (Section 6.3.2). We will show that each arc cost can be computed in $\mathcal{O}(T^2)$ time (Section 6.3.3). Altogether, this leads to an $\mathcal{O}(T^7/\ell)$ time algorithm for $(\mathcal{P}^{(\ell)}(\hat{b}))$ when both setup and production expenditures are time-invariant. In the following, we present some structural properties that will be used in Sections 6.3.2 and 6.3.3.

6.3.1. Structural Properties

In this section, we prove some structural properties that will allow us to find the production quantities in a maximal sequence of subplans. The following lemmas will be used to prove that, except for the first one, all blocks fully contained in the maximal sequence are extreme. We say that a block is split if it contains at least two (partial) subplans. For example, in Fig. 1, Block 3 is a split block. Furthermore, the block containing the period v is denoted by $[b(v), e(v)]$.

LEMMA 6.3: Each production period in a subplan belongs to a different block.

LEMMA 6.4: There exists an optimal solution for which, except for the first one, all production periods in a subplan are contained in a tight block.

LEMMA 6.5: Let $[u, v]$ and $[v + 1, w]$ be connected subplans starting in different blocks. Then, the connecting block, $[b(v), e(v)]$, is tight.

PROPOSITION 6.6: Let $[t, w]$ be a maximal sequence of connected subplans. Except for the first one, all blocks fully contained in the maximal sequence are extreme. Moreover, if the last block of the sequence $[b(w), e(w)]$ is split, then there is no production in $[b(w), w]$.

As a consequence of Proposition 6.6, we are now able to determine the relevant production quantities, given the number of tight blocks and setups in the remainder of the maximal sequence. This is formalized in Propositions 6.8 and 6.9, where the specific production quantities are given by the following definition.

DEFINITION 6.7: Define $s_{v,w}^{\eta,\gamma} = d_{v,w} - \eta\tilde{b} + \gamma\tilde{f}$ and $r_{v,w}^{\eta,\gamma} = \eta\tilde{b} - \gamma\tilde{f} - d_{v,w}$.

PROPOSITION 6.8: Let $[t, w]$ be a maximal sequence of connected subplans. Let v be the last regeneration period of a block in the sequence with $v < w$. Let η be the total number of tight blocks in $[e(v) + 1, w]$ and γ be the total number of setups placed in $[e(v) + 1, w]$. Then $x_{v+1} = s_{v+1,w}^{\eta,\gamma}$.

PROPOSITION 6.9: Let $[t, w]$ be a maximal sequence of connected subplans. Let $v < w$ be the first regeneration period of a split block. Let η be the total number of tight blocks in $[e(v) + 1, w]$ and γ the total number of setups placed in $[v+1, w]$. The only production in $[b(v), v]$ is at most $r_{v+1,w}^{\eta+1,\gamma+1}$.

We illustrate these results using the maximal sequence [6, 16] in Fig. 1. From Proposition 6.6, Block 4 must be tight, while from the figure we see that we incur there one single setup. For Block 3, the first and last regeneration points are 9 and 11, respectively. Proposition 6.8 implies that $x_{12} = s_{12,16}^{1,1}$, while from Proposition 6.9 $x_9 = r_{10,16}^{2,4}$.

6.3.2. The Inner DP Approach

As aforementioned, to find the optimal solution of a maximal sequence $[t, w]$, each maximal sequence is split into a collection of subplans. In the following, we will show that we can use a network with $\mathcal{O}(T^4/\ell)$ arcs to construct a maximal sequence with minimum cost.

We will assume without loss of generality that $t = b(t)$ and $w = e(w)$. If $t > b(t)$, using Proposition 6.6, there is no production in $[b(t), t - 1]$. We can solve instead a new sequence, $[t', w]$, beginning at period $t' = b(t)$, and where the demand between t' and $t - 1$ is equal to zero, while there is no production in $[t', t - 1]$. If $w < e(w)$, again using Proposition 6.6, there is no production in $[b(w), w]$. We can solve instead a new sequence, $[t, w']$, ending at period $w' = b(w) - 1$,

and where the demand at the new end period w' is equal to $d_{b(w)-1,w}$.

To find the optimal cost of a maximal sequence $[t, w]$, we use a shortest path approach. A node (u, η, γ) represents a partial solution for periods $[u, w]$ with $u - 1$ a regeneration period, η the number of tight blocks in $[e(u) + 1, w]$, and γ the number of setups in $[u, w]$. Since we construct the maximal sequence moving backwards in time, the initial node is $(w + 1, 0, 0)$. There are two types of arcs, representing either a subplan spanning across multiple blocks or a subplan contained in a block.

Let us discuss the first type of arcs, that is, a subplan $[u, v]$ spanning across multiple blocks. In that case, $u < v$ will be two periods in different blocks, that is, $e(u) < e(v)$. Given a partial solution represented by $(v + 1, \eta, \gamma)$, the demand $d_{u,v}$ uniquely determines the number of tight blocks and the number of setups in $[u, v]$, say η^{uv} and γ^{uv} , respectively, as shown in the appendix. Hence, we have an arc from $(v + 1, \eta, \gamma)$ to $(u, \eta + \eta^{uv} + 1, \gamma + \gamma^{uv})$. Note that we have $\eta + \eta^{uv} + 1$ tight blocks, since the connecting block $[b(v), e(v)]$ is tight. Clearly, for a given node $(v + 1, \eta, \gamma)$, the number of outgoing arcs of this type is $\mathcal{O}(T)$.

We now present the second type of arcs, that is, a subplan $[u, v]$ contained in a block. In that case, $u \leq v$ will be two periods in the same block, and we have an arc from $(v + 1, \eta, \gamma)$ to $(u, \eta, \gamma + 1)$. Clearly, for a given node $(v + 1, \eta, \gamma)$, the number of outgoing arcs is $\mathcal{O}(\ell)$. Finally, to complete the network, we add a sink node with incoming arcs from the nodes (t, η, γ) with costs zero.

We now count the total number of arcs in the network corresponding to a maximal sequence $[t, w]$. The total number of nodes (u, η, γ) is $\mathcal{O}(T \cdot \frac{T}{\ell} \cdot T) = \mathcal{O}(T^3/\ell)$. Since the number of incoming type 1 and type 2 arcs is $\mathcal{O}(T)$ and $\mathcal{O}(\ell)$, the total number of type 1 and type 2 arcs is $\mathcal{O}(T^4/\ell)$ and $\mathcal{O}(T^3)$, respectively. Hence, the optimal cost of a maximal sequence $[t, w]$ can be found by solving a shortest path problem on a network with $\mathcal{O}(T^4/\ell)$ arcs.

Note that when solving the whole problem, $(\mathcal{P}^{(\ell)}(\hat{b}))$, we can achieve some savings. Since the number of maximal sequences is $\mathcal{O}(T^2)$, a straightforward approach leads to a total number of $\mathcal{O}(T^6/\ell)$ type 1 arcs and $\mathcal{O}(T^5)$ type 2 arcs. However, the arcs used in the computation of the maximal sequence $[t, w]$ are a subset of the arcs used in the maximal sequence $[t - 1, w]$. Hence, the total number of different type 1 and type 2 arcs is equal to $\mathcal{O}(T^5/\ell)$ and $\mathcal{O}(T^4)$, respectively.

6.3.3. Determining the Optimal Cost of a Subplan

Subplans Spanning Across Multiple Blocks. Consider a subplan $[u, v]$ spanning across multiple blocks which is not the last one in a maximal sequence, and hence the predecessor of the partial solution $(v + 1, \eta, \gamma)$. The case where it is the last

one can be handled in a similar way and hence is omitted. We will sketch how to develop a DP approach yielding the optimal cost for such a subplan. The details can be found in the appendix.

From the results derived so far, we can determine the production amounts. In the appendix, we show that the last and first production quantities of the subplan can be determined by Proposition 6.8 and Proposition 6.9, respectively. Furthermore, since each production period in a subplan belongs to a different block (Lemma 6.3) and this block is tight (Lemma 6.4), any other production amount equals $\tilde{b} - \hat{f}$.

Hence, what remains is the allocation of these production amounts (if possible). We propose an extension to the approach in [12], where we allocate the full production periods to blocks, and then to a period within the chosen block. In the appendix, we show how this can be done in $\mathcal{O}(T^2)$ time for a single subplan. Since the total number of type 1 arcs is $\mathcal{O}(T^5/\ell)$ (see Section 6.3.2), the total running time spent on computing type 1 arc costs is $\mathcal{O}(T^7/\ell)$.

Subplans Contained in a Block. Consider a subplan covering the interval $[u, v]$, contained in block $[b(u), e(u)] = [b(v+1), e(v+1)]$, with η tight blocks in $[e(v+1)+1, w]$ and γ setups in $[v+1, w]$, represented by an arc from $(v+1, \eta, \gamma)$ to $(u, \eta, \gamma+1)$. In this case, checking feasibility of a single arc is a trivial exercise and takes constant time. As usual, the cost is set to infinity in case of infeasibility. We know that the number of setups placed in $[u, e(u)]$ is equal to $\gamma+1$. By Proposition 6.9, the remaining capacity is $r_{v+1,w}^{\eta+1,\gamma+1}$, and hence the subplan is feasible if $d_{u,v} \leq r_{v+1,w}^{\eta+1,\gamma+1}$. As an example, subplan [10, 11] in Fig. 1, contained in Block 3, will be feasible if $d_{10,11} \leq r_{12,16}^{2,3}$.

The cost of a single arc can be found in constant time by noting that the cost of subplan $[u, v+1]$ can be computed from subplan $[u, v]$ (similar to Section 6.2). Hence, checking feasibility and computing the cost can be done in constant time for a single arc. Since the total number of type 2 arcs is $\mathcal{O}(T^4)$ (see Section 6.3.2), the total running time spent on type 2 arcs is $\mathcal{O}(T^4)$.

6.4. Time-Invariant Production Expenditures

In this section, we deal with the case where the lot-sizing costs are such that there are no speculative motives to hold inventory, $\hat{f}_t = \hat{h}_t = 0$, and \hat{c}_t are time-invariant. The analysis is similar to the one of the previous section. The main difference is that the state space of the inner DP algorithm can be reduced. Since there are no setup expenditures, we do not have to keep track of the number of setups. Moreover, the total demand in an interval $[u, w]$ consisting of connected subplans exactly determines the number of tight blocks, namely $\lfloor d_{u,w}/\hat{b} \rfloor$. Therefore, to find the optimal cost of a maximal sequence $[u, w]$, it is sufficient to define a network with nodes

Table 3. Complexity overview in case of block expenditures.

f_t	Costs		Expenditures			Complexity
	c_t and h_t	\hat{f}_t	\hat{c}_t	\hat{h}_t		
v	Nonspeculative	c	c	0	Polynomially solvable	
c	Nonspeculative	v	0	0	\mathcal{NP} -hard	
c	Nonspeculative	0	v	0	\mathcal{NP} -hard	
c	Nonspeculative	0	0	c	\mathcal{NP} -hard	

(v), a partial solution for periods $[v, w]$ where $v-1$ is a regeneration period. It can be verified that the number of nodes reduces from $\mathcal{O}(T^3/\ell)$ to $\mathcal{O}(T)$, and the number of arcs from $\mathcal{O}(T^5/\ell)$ to $\mathcal{O}(T^3)$. Since the running time for the construction of the arcs remains unchanged and hence equals $\mathcal{O}(T^2)$, the total running time becomes $\mathcal{O}(T^5)$.

6.5. Discussion of the Complexity Results

As aforementioned, any class of instances that is \mathcal{NP} -hard for either the whole horizon or the period models is also \mathcal{NP} -hard for the block model. An overview of the obtained complexity results is given in Table 3. We only show the classes of \mathcal{NP} -hard instances with the strongest assumptions on the parameters, which turn out to be the \mathcal{NP} -hard instances for the whole horizon model. As for the whole horizon model, we see that relaxing any of the assumptions on the polynomially solvable instances leads to an \mathcal{NP} -hard problem. Therefore, we can conclude that the complexity results for the block model are tight as well.

7. ON THE PARETO FRONTIER

The Pareto frontier of $(\text{BOLS}^{(\ell)})$ can be used to describe the tradeoff between lot-sizing costs and expenditure and whether expenditure has a great impact on lot-sizing costs. In this section, we examine the shape of the Pareto frontier and the task of describing it. In Section 7.1, we study a polynomially solvable case of the whole horizon model $(\text{BOLS}^{(T)})$. In particular, under the class of instances in Section 4.2, with time-invariant f_t, c_t, \hat{f}_t , and \hat{c}_t , and $\hat{h}_t = \alpha h_t$ for some $\alpha \geq 0$, we show that the Pareto frontier has $\mathcal{O}(T)$ points and is convex, and it can be described in $\mathcal{O}(T^2)$ time. In Section 7.2, we illustrate that none of these properties will hold in general. For $(\text{BOLS}^{(T)})$, we illustrate this with an instance where the parameters are time-variant, while for $(\text{BOLS}^{(1)})$, we use one where $(\mathcal{P}^{(1)}(\hat{b}))$ is polynomially solvable, described in Section 5.2. In Section 7.3, we outline how to approximate a subset of the Pareto frontier of $(\text{BOLS}^{(\ell)})$, $\ell = 1, \dots, T-1$, by using the results of Sections 5.2 and 6.3 on $(\mathcal{P}^{(\ell)}(\hat{b}))$.

7.1. A Polynomially Solvable Case for (BOLS^(T))

In this section, we show that for the class of instances in Section 4.2, where the production parameters are time-invariant and the inventory holding parameters satisfy $\hat{h}_t = \alpha h_t$, the Pareto frontier has $\mathcal{O}(T)$ points, is convex, and can be described in $\mathcal{O}(T^2)$ time.

PROPOSITION 7.1: If f_t, c_t, \hat{f}_t , and \hat{c}_t are time-invariant, and $\hat{h}_t = \alpha h_t$ for some $\alpha \geq 0$, the Pareto frontier of (BOLS^(T)) satisfies the following properties

1. it has $\mathcal{O}(T)$ points,
2. it is convex, and
3. it can be described in $\mathcal{O}(T^2)$ time.

PROOF: Recall from Section 4.2 that, for this class of instances, once the number of setups n is fixed, minimizing the lot-sizing costs is equivalent to minimizing the expenditure across the whole planning horizon. Given n , let $(\mathcal{L}_n, \mathcal{E}_n)$, the lot-sizing costs and expenditure of the solution with minimum lot-sizing costs among those having n setups. It is easy to show that $(\mathcal{L}_n, \mathcal{E}_n)$ is the only possible Pareto efficient outcome with n setups. Thus, the first claim follows.

To show the convexity of the Pareto frontier, it is sufficient to prove that there are no other Pareto efficient outcomes than the supported ones, that is, Pareto efficient outcomes on the lower convex envelope of the Pareto frontier. It is well-known that the supported outcomes can be found by taking a convex combination of the lot-sizing costs and the expenditure. Given $\lambda \in [0, 1]$, finding the corresponding supported outcome boils down to solving an ELS problem with the following objective function

$$\sum_{t=1}^T [\lambda(fy_t + cx_t + h_t I_t) + (1 - \lambda)(\hat{f}y_t + \hat{c}x_t + \hat{h}_t I_t)],$$

which, using $\hat{h}_t = \alpha h_t$, can be rewritten as

$$\sum_{t=1}^T [(\lambda f + (1 - \lambda)\hat{f})y_t + (\lambda c + (1 - \lambda)\hat{c})x_t + (\lambda + (1 - \lambda)\alpha)h_t I_t].$$

After ignoring the variable production cost term (it is just a constant) and scaling, we obtain

$$\sum_{t=1}^T \left[\left(\frac{\lambda}{\lambda + (1 - \lambda)\alpha} f + \frac{(1 - \lambda)}{\lambda + (1 - \lambda)\alpha} \hat{f} \right) y_t + h_t I_t \right].$$

Thus, under this class of instances, finding all supported outcomes of (BOLS^(T)) can be viewed as performing a parametric analysis on the setup cost of an ELS problem with nonspeculative motives. This parametric problem has been

studied by [24] in the case that the setup cost changes in an additive way, that is considering setup cost changes of the form $f_t - \delta$, in contrast to a multiplicative way, that is considering setup cost changes of the form δf_t . However, in case of time-invariant setup cost, there is a one-to-one correspondence between both additive and multiplicative changes, and hence we can use their results. In particular, Ref. [24] shows that when performing a parametric analysis, the number of setups in an optimal solution change one-by-one in a structured way, and the value of λ for which the number of setups changes from n to $n + 1$ can be found in linear time. In particular, it follows from their analysis that no solution with n setups is “skipped over,” that is, if $(\mathcal{L}_n, \mathcal{E}_n)$ (resp. $(\mathcal{L}_{n+2}, \mathcal{E}_{n+2})$) is a Pareto efficient outcome with n (resp. $n + 2$) setups in the frontier, we also find a Pareto efficient outcome $(\mathcal{L}_{n+1}, \mathcal{E}_{n+1})$ with $n + 1$ setups. (For the details, we refer to [24]; see the proof of Theorem 4.) Combining this with claim 1 of this proposition, the desired result follows.

To construct the Pareto frontier, we use the following approach. We start with $\lambda = 0$ and add the corresponding outcome to the Pareto frontier. Because the value of λ for which the number of setups changes can be found in linear time (as mentioned in the previous part), the next Pareto efficient outcome can be found in $\mathcal{O}(T)$ time. In turn, because the frontier consists of $\mathcal{O}(T)$ points, it can be constructed in $\mathcal{O}(T^2)$ time. □

7.2. The Shape of the Pareto Frontier of (BOLS^(ℓ))

These two well-behaving properties of the Pareto frontier, convexity, and polynomiality of the number of points, do not hold in general. In this section, we present counterexamples for (BOLS^(T)) and (BOLS⁽¹⁾). Similar conclusions can be drawn for the block case, (BOLS^(ℓ)), using the fact that it is a generalization of the period model.

The 15-period instance with cost, expenditure, and demand parameters given in Table 4 illustrates that the Pareto frontier of (BOLS^(T)), shown in Fig. 2, has a point, (877, 50), which is not in the convex envelope of the Pareto efficient outcomes. The figure has been constructed by solving $(\mathcal{P}^{(T)}(\hat{b}))$ for different values of \hat{b} (note that multiples of 5 are sufficient to find all strongly Pareto efficient outcomes, because we only have setup expenditures, and those are multiples of 5). Thus, the Pareto frontier is not convex for this instance, where the production parameters are time variant. A nonconvex Pareto frontier is also observed in [6] in an economic order quantity setting.

In addition, there are instances for which the number of points in the Pareto frontier of (BOLS^(T)) is exponential. Consider again the $2n$ -period instance from Proposition 4.1 given in Table 5. Let $S \subset \{1, \dots, n\}$ and consider a solution where the production periods are given by $t = 2i - 1, i \in S$, and $t = 2i, i \in S^c$. It follows from the proof of Proposition 4.1 that

Table 4. The data for the 15-period instance of the whole horizon model.

c_t	h_t	\hat{c}_t	\hat{h}_t												
5	1	0	0												
t	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
f_t	25	25	25	25	25	25	25	10	25	25	25	25	25	25	25
\hat{f}_t	10	10	10	10	10	10	10	25	10	10	10	10	10	10	10
d_t	10	10	4	2	1	40	10	10	10	10	10	3	4	7	3




Figure 2. The Pareto frontier of (BOLS^(T)) for the 15-period instance in Table 4. [Color figure can be viewed in the online issue, which is available at wileyonlinelibrary.com.]

Table 5. Instance from Proposition 4.1.

t	d_t	f_t	c_t	h_t	\hat{f}_t	\hat{c}_t	\hat{h}_t
$2i - 1$	0	A	0	a_i	0	0	0
$2i$	1	A	0	∞	0	a_i	0

the total costs equal $nA + \sum_{i \in S} a_i$, while the total expenditure equals $\sum_{i \in S^c} a_i = 2A - \sum_{i \in S} a_i$. Therefore, each Pareto efficient outcome is of the form $(nA + a, 2A - a)$ for some $a \in \mathbb{N}$. It follows that checking whether the point $(nA + a, 2A - a)$ belongs to the Pareto frontier boils down to the question whether there exists a subset S such that $\sum_{i \in S} a_i = a$. The latter problem is the well-known \mathcal{NP} -complete subset sum problem. Moreover, if the values a_i have enough variation (e.g., if $\sum_{i \in S} a_i \neq \sum_{i \in S'} a_i$ for $S \neq S'$), then the number of Pareto efficient solutions is of exponential order.

We now turn to the period model (BOLS⁽¹⁾) and show that the Pareto frontier may have nonconvex sections, even for instances where $(\mathcal{P}^{(1)}(\hat{b}))$ is polynomially solvable. Consider the 15-period instance with parameters given in Table 6, where all production and inventory parameters are stationary. Its Pareto frontier, shown in Fig. 3, has nonconvex sections (e.g., when the expenditure values range between 150 and 200).

The shape of the Pareto frontier seen in Figs. 2 and 3 can be explained using arguments that hold for any biobjective

mixed integer problem. The Pareto frontier can be partitioned into maximal intervals $[\hat{b}_L, \hat{b}_U]$ on which it is piecewise linear and convex, as in Fig. 3, and the y -vector is the same. To see why this holds, consider (BOLS^(ℓ)) in which we fix the y -vector. Finding the Pareto frontier of this restricted problem is equivalent to performing a parametric analysis on \hat{b} of $(\mathcal{P}^{(\ell)}(\hat{b}))$, which is now a linear programming (LP) problem. The value \hat{b}_L is the point where $(\mathcal{P}^{(\ell)}(\hat{b}))$ under y becomes infeasible, while \hat{b}_U is the point where the expenditure constraints are not binding anymore. From LP theory, we know that the result of this parametric analysis is a piecewise linear and convex function on $[\hat{b}_L, \hat{b}_U]$. Now the result follows by noting that the Pareto frontier is equal to the lower envelope of a finite number of piecewise linear and convex functions (namely, for each feasible solution of y -variables we have such a function). Note that if the frontier consists of points only such as in Fig. 2, then we are dealing with the special case where $\hat{b}_L = \hat{b}_U$.

7.3. Constructing ϵ -Dominating Sets of the Pareto Frontier

In this section, we focus on the instances of (BOLS^(ℓ)), $\ell = 1, \dots, T - 1$, for which $(\mathcal{P}^{(\ell)}(\hat{b}))$ is polynomially solvable. It is not difficult to verify that using a weighted approach as in Section 7.1, does not lead to a problem which is easier than $(\mathcal{P}^{(\ell)}(\hat{b}))$. Therefore, even finding supported solutions of (BOLS^(ℓ)) seems not an easy task. Therefore, instead of trying to construct the supported points of the frontier, one can make use of the concept of ϵ -dominating sets (see [5] and references therein) to approximate a subset of the Pareto frontier in a running time which is pseudopolynomial in the input size and polynomial in the inverse of the required precision $\frac{1}{\epsilon}$ assuming $(\mathcal{P}^{(\ell)}(\hat{b}))$ is polynomially solvable. (Note that we exclude $\ell = T$ since in this case the Pareto frontier can be described in polynomial time as shown in Proposition 7.1.)

To formalize the concept ϵ -dominating set, consider the following general biobjective problem

$$\min_{z \in Z} (f(z), g(z)).$$

Table 6. The data for the 15-period instance of the period model.

f_t	c_t	h_t	\hat{f}_t	\hat{c}_t	\hat{h}_t										
25	5	1	0	5	1										
t	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
d_t	10	10	4	2	1	40	10	10	10	10	10	3	4	7	3




Figure 3. The Pareto frontier of (BOLS⁽¹⁾) for the 15-period instance in Table 6. [Color figure can be viewed in the online issue, which is available at wileyonlinelibrary.com.]

DEFINITION 7.2: The set Z^* is called an ε -dominating set in the value space if, for each $z \in Z$, there exists a $z^* \in Z^*$ such that

$$f(z) \geq f(z^*) - \varepsilon \quad \text{and} \\ g(z) \geq g(z^*) - \varepsilon.$$

Given upper and lower bounds on \hat{b} , the following straightforward algorithm can be used to construct an ε -dominating set for (BOLS^(ℓ)).

Algorithm for finding an ε -dominating set on $[L, U]$

- Step 0.** Set $Z^* = \emptyset$. Let L and U be a lower bound and an upper bound on \hat{b} . Let $\{\hat{b}^i\}$ be a grid of $[L, U]$, such that $\hat{b}^{i+1} - \hat{b}^i = \varepsilon$.
- Step 1.** For each i , solve $(\mathcal{P}^{(\ell)}(\hat{b}^i))$ and add its optimal solution to Z^* .

For the polynomially solvable instances in Sections 5–6, it is trivial to see that this algorithm runs in polynomial time in $U - L$, $1/\varepsilon$ and T . Clearly, if the interval $[L, U]$ is relatively small, then the above algorithm will be able to find the ε -dominating set in a reasonable amount of computation

time. However, the running time is pseudopolynomial in general, because the number $U - L$ may be pseudopolynomial in the input size of the instance. Clearly, for fixed L and U , for example, specified by a decision maker, the running time of this algorithm becomes polynomial.

8. CONCLUSIONS

In this article, we study a biobjective Economic Lot-Sizing model, (BOLS^(ℓ)), which arises, when not only the lot-sizing costs, but also expenditures are a concern. The parameter ℓ defines the level of aggregation used to measure the expenditure, the lower ℓ the more granular we are in terms of recording the expenditure. Apart from the truly block case, the two extreme cases are also relevant. In the whole horizon case, (BOLS^(T)), the level of aggregation is the same for both the lot-sizing costs and the expenditure, while in the period case, (BOLS⁽¹⁾), we record the expenditure for each period.

Besides incorporating environmental issues into the lot-sizing problem, (BOLS^(ℓ)) can also be used when restricting the type of plan that the classical lot-sizing model may return. With the modelling in this article, we can impose a constraint in each block to have a plan that spreads out the burden, say the number of setups, amount of inventory, or even the total lot-sizing costs.

We have shown that the Pareto efficient outcome problem, $(\mathcal{P}^{(\ell)}(\hat{b}))$, is \mathcal{NP} -hard in general, and we have identified non-trivial problem classes for which this problem is polynomially solvable. Being able to solve $(\mathcal{P}^{(\ell)}(\hat{b}))$ in polynomial time is not only important on its own, but it is also relevant when describing segments of the efficient frontier defined by the decision maker.

We have shown that $(\mathcal{P}^{(T)}(\hat{b}))$ can be solved in $\mathcal{O}(T^2)$ time in the almost time-invariant case, in which all production parameters are time-invariant while for the inventory holding parameters we assume $\hat{h}_t = \alpha h_t$. For the period case, we have shown that $(\mathcal{P}^{(1)}(\hat{b}))$ can be solved in $\mathcal{O}(T^2)$ time if the lot-sizing costs are such that there are no speculative motives to hold inventory, the setup costs are nondecreasing and the expenditure parameters are time-invariant. For the general block case, we have shown that $(\mathcal{P}^{(\ell)}(\hat{b}))$ can be solved in $\mathcal{O}(T^7/\ell)$ time if the lot-sizing costs are such that there are no speculative motives to hold inventory, the

production expenditure parameters are time-invariant, and there are no inventory expenditures, where the running time can be improved to $\mathcal{O}(T^2\ell)$ and $\mathcal{O}(T^5)$ for the corresponding single-process expenditure cases. These polynomiality results are tight for the whole horizon and the block cases, while for the period case, they are tight, except for one single case. Indeed, it remains an open question whether $(\mathcal{P}^{(1)}(\hat{b}))$ under inventory expenditures only is an \mathcal{NP} -hard problem.

In addition to the complexity of the open problem and the issues mentioned at the end of Section 3, there are other interesting lines of future research. The first one is the study of generalizations of the problem, for example, problems with general cost parameters and production capacities. The second one is an alternative way of aggregating the expenditures. In $(\text{BOLS}^{(\ell)})$, the blocks in which the expenditure are recorded define a partition of the planning horizon. We are currently investigating the case in which the expenditures are recorded in every block of ℓ consecutive periods, and therefore the blocks may overlap. The third one is the development of solution approaches for the \mathcal{NP} -hard instances of $(\mathcal{P}^{(\ell)}(\hat{b}))$, based on strong mixed integer programming formulations, finding valid inequalities, or developing FPTASes. The fourth one is the approximation of the Pareto frontier itself by for example a constant factor approximation or an FPTAS (instead of a pseudopolynomial algorithm), see also [17] and [18]. The fifth one is the development of exact approaches to describe the entire Pareto frontier for some classes of the problem, see for example [20].

APPENDIX

In the following, we present the proofs of the propositions and lemmas in Section 6 as well as an $\mathcal{O}(T^2)$ time algorithm for the optimal cost of a subplan spanning across multiple blocks.

PROOF OF PROPOSITION 6.1: Consider an optimal solution for which the ZIO property does not hold, that is, there exists a period t with $I_{t-1}x_t > 0$. Let s be the last production period before period t (this period exists since $I_0 = 0$). By decreasing production in period s and increasing production in period t by the same sufficiently small amount, the solution remains feasible w.r.t. the inventory balance constraints. Furthermore, the inventory levels in periods $[s, t-1]$ will decrease, and there are no additional setups in the modified solution. This means that (i) the lot-sizing costs will not increase because of the nonspeculative motives assumption, and (ii) the solution remains feasible w.r.t. the expenditure constraints. It is not difficult to see that by choosing the change in production amount as large as possible, and by repeating the above procedure, we end up with a solution satisfying the ZIO property and having equal or lower lot-sizing costs, and the desired result follows. \square

PROOF OF LEMMA 6.3: The result follows since the ZIO property holds within blocks.

Indeed, let u and v be two consecutive production periods within a block. We will show below that there exists an optimal extreme point solution with $I_{v-1} = 0$. Since the production expenditures are time-invariant, moving production within a block does not effect the expenditures. Because the costs are nonspeculative, the total costs will not increase by postponing production in

some period. By postponing, the highest amount in period u while keeping feasibility, we obtain a solution with $I_{v-1} = 0$. \square

PROOF OF LEMMA 6.4: Recall that each production period in a subplan belongs to a different block, see Lemma 6.3. Therefore, it is enough to show that if we have two consecutive production blocks with production in the subplan, say blocks i and j , then block j should be tight. The result follows using the nonspeculative motives of the lot-sizing costs.

Suppose that block j is nontight, and therefore it has spare capacity $\delta_j > 0$, which can be used to increase the production level in any of the existing production periods in the block. Let t be the only production period from block j in the subplan and t' be the last production period before t . Therefore, period t' belongs to block i and also to the subplan. Recall that within a subplan all the inventory levels are strictly positive. We can reduce the production in period t' as well as the inventory levels in periods $t', t'+1, \dots, t-1$ by ε , at the same time that we increase the production in period t by ε .

Constraint (8) is still satisfied for block i since we have reduced the corresponding left hand side, and perhaps increased the right hand side if $\varepsilon = x_{t'}$. Constraint (8) for block j is still satisfied if $\varepsilon \leq \delta_j$. Furthermore, the production in period t' cannot be negative, which means that $\varepsilon \leq x_{t'}$. Finally, we also need to impose that $\varepsilon \leq I_{t-1}$ to make sure that the new inventory levels in $t', t'+1, \dots, t-1$ are all nonnegative. Thus, for an appropriate choice of ε , namely $\varepsilon = \min\{x_{t'}, I_{t-1}, \delta_j\}$, this solution is still feasible. Observe that because of the nonspeculative motives assumption of the lot-sizing costs, this solution is also optimal.

If $\varepsilon = x_{t'}$, then the result holds for blocks i and j since t' is not a production period anymore. If $\varepsilon = I_{t-1}$, then the new inventory level at the end of period $t-1$ is equal to zero and the subplan decomposes into two new subplans. Finally, if $\varepsilon = \delta_j$, then block j becomes tight, and the desired result holds again for blocks i and j . \square

PROOF OF LEMMA 6.5: Note that $v+1 \in [b(v), e(v)]$ because the subplans are connected, while by definition $u \in [b(u), e(u)]$. Because subplans $[u, v]$ and $[v+1, w]$ start in different blocks, this means that block $[b(v), e(v)]$ is different from $[b(u), e(u)]$. Therefore, block $[b(v), e(v)]$ is not the first block of the subplan $[u, v]$, and by Lemma 6.4, it must be tight. \square

PROOF OF PROPOSITION 6.6: The first part of the result follows from the fact that (i) a block fully contained in a subplan is extreme (see Lemma 6.4), and (ii) a block associated with two connected subplans starting in different blocks is extreme as well (see Lemma 6.5). Since a block fully contained in $[t, w]$ (which is not the first one) is one of such blocks, it must be extreme.

The second part of the result is trivial. If the last block $[b(w), e(w)]$ is split, that is, $w < e(w)$, there should not be production in $[b(w), w]$, since otherwise the sequence would not end at w . \square

PROOF OF PROPOSITION 6.8: Since v is the last regeneration period in its block and $v < w$, we have that $v < e(v)$ and hence the interval $[v+1, e(v)]$ is well-defined. By Proposition 6.6, all blocks following block $[b(v), e(v)]$ (if any) are extreme (in case of a complete block) or have no production (in case the last block is split). With the definition of η and γ , we have that the total production in $[e(v)+1, w]$ is equal to $\eta\bar{b} - \gamma\bar{f}$. Therefore, the total production in $[v+1, e(v)]$ equals $d_{v+1,w} - (\eta\bar{b} - \gamma\bar{f})$. Since there is only a single production in $[v+1, e(v)]$ which must be in period $v+1$, the desired result follows. \square

PROOF OF PROPOSITION 6.9: First, we will show that there is exactly one production in $[b(v), v]$. Since $v < w$, and because of the connectivity, we know that there should be production in $[b(v), v]$. Now, the uniqueness follows from the fact that v is the first regeneration period of block $[b(v), e(v)]$.

Note that the amount of production in $[b(v), v]$, say x , is maximal if it makes the block $[b(v), e(v)]$ tight. So suppose that the block is tight. Then, the total expenditure capacity of the tight blocks in $[v, w]$ should be equal to the total expenditure in $[v, w]$. The first is equal to $(\eta + 1)\tilde{b}$. Moreover, the latter is equal to the expenditure resulting from the total amount of units produced $(x + d_{v+1,w})$ plus the expenditure resulting from the setups used $((\gamma + 1)\hat{f})$. Equating these amounts gives $x = (\eta + 1)\tilde{b} - (\gamma + 1)\hat{f} - d_{v+1,w} = r_{v+1,w}^{\eta+1,\gamma+1}$. \square

Determining Optimal Subplans Spanning Across Multiple Blocks in $\mathcal{O}(T^2)$ Time

Consider a subplan $[u, v]$ spanning across multiple blocks which is not the last one in a maximal sequence, and hence the predecessor of the partial solution $(v + 1, \eta, \gamma)$. In this section, we develop a DP approach yielding the optimal cost for such a subplan. In case the subplan is the last one in a maximal sequence, the required modifications are straightforward.

Let m be the number of production periods in the subplan. We will show that the subplan starts and ends with some specific production quantity and has $m - 2$ full production periods, that is, where $\tilde{b} - \tilde{f}$ units are produced, in between. The optimal cost will then be calculated using an approach similar to [12]. Note that the parameters η^{uv} and γ^{uv} in Section 6.3.2 equal $m - 2$ and m , respectively. For ease of notation, we use the single parameter m in this section.

Since subplan $[u, v]$ is an intermediate subplan or the first one in the maximal sequence, we know that the subplan will have a single production period in the interval $[b(v), v]$. Moreover, it follows from Proposition 6.9 and the fact that the block is tight, that the production quantity is equal to $r_{v+1,w}^{\eta+1,\gamma+1}$. Using Lemma 6.3, Lemma 6.4 and the value of the fractional productions, we derive the following corollary.

COROLLARY A.1: We have that $m = \left\lfloor \frac{d_{u,v} - r_{v+1,w}^{\eta+1,\gamma+1}}{\tilde{b} - \tilde{f}} \right\rfloor + 2$.

We can now also identify the production quantity in period u . It follows that there are $\eta + m - 1$ tight blocks and $\gamma + m - 1$ setups in $[e(u + 1), w]$. Hence, by applying Proposition 6.8 to $[u, e(u)]$, we should have $x_u = s_{u,w}^{\eta+m-1,\gamma+m-1}$. (Note that if $\tilde{b} - \tilde{f}$ divides $d_{u,v} - r_{v+1,w}^{\eta+1,\gamma+1}$, then we have a full production in period u and the expenditure constraint is tight for periods $[u, e(u)]$. This is only possible in case $[u, v]$ is the first subplan in the sequence. Otherwise, the subplan cannot be part of a maximal sequence.)

To develop the DP algorithm, we introduce some notation to define the range of feasible periods for each production. Let $n_b = (e(v) - e(u)) / \ell$ be the number of blocks covered by the subplan (including the blocks $[b(u), e(u)]$ and $[b(v), e(v)]$). Let $n(t) \in \{1, \dots, n_b\}$ be the block number of period t in the interval $[u, v]$, where we start counting at the block that contains period u , so $n(u) = 1$. Let t_i be the latest feasible period for the i th production period in the subplan. So,

$$t_i = \begin{cases} u & i = 1 \\ \max\{j : d_{u,j-1} < x_u + (i-1)(\tilde{b} - \tilde{f})\} & i = 2, \dots, m-1 \\ \max\{j : d_{u,j-1} < x_u + (m-2)(\tilde{b} - \tilde{f}) + r_{v+1,w}^{\eta+1,\gamma+1}\} & i = m. \end{cases}$$

To have a feasible solution, we should have $n(t_i) \geq i$, since we need at least i blocks for the first i production periods, and $t_m \geq b(v)$, to make sure that the last production occurs in $[b(v), v]$. Note that, if these inequalities are satisfied, it is feasible to associate a production period with every period t_i . Otherwise, we set the cost of the subplan to infinity to reflect its infeasibility.

The following proposition identifies the range of feasible periods for each full production. (Note that we are assuming that the inventory holding costs have been incorporated into the production costs using the balance constraints, allowing us to talk about the cheapest full production period.)

PROPOSITION A.2: The optimal production period associated with the i th production is either the cheapest full production period in a complete block before block number $n(t_i)$, or the cheapest full production period in the partial block $[b(t_i), t_i]$.

PROOF: In the first case, every period in the full block is feasible by construction of t_i and hence we can choose the cheapest period in the block. In the second case, only periods in $[b(t_i), t_i]$ are feasible and hence we choose the cheapest among these. \square

Using this proposition, we are now ready to present the DP algorithm. We use the recursion variable $f(i, j)$ for $i = 1, \dots, m$ and $j = i, \dots, n(t_i)$: the minimum cost of assigning i production periods in $[u, t_i]$ when only using blocks $1, \dots, j$. In case $j = n(t_i)$, the block j is the partial block $[b(t_i), t_i]$. Note that the first and the last production quantities equal $s_{u,w}^{\eta+m-1,\gamma+m-1}$ and $r_{v+1,w}^{\eta+1,\gamma+1}$, respectively, while the other quantities are equal to $\tilde{b} - \tilde{f}$. Using the DP variable $f(i, j)$ and letting $f(0, 0) = 0$ and $f(1, 0) = \infty$, the following recursions hold:

$$f(i, j) = \min\{f(i, j-1), f(i-1, j-1) + c(j)\},$$

where $c(j)$ is the lowest cost of a full production period in block number j (again the block j may be split), and

$$f(m, n_b) = f(m-1, n_b-1) + c(n_b),$$

where, with some abuse of notation, $c(n_b)$ is the cost of producing $r_{v+1,w}^{\eta+1,\gamma+1}$ in the cheapest period of $[b(v), t_m]$. Clearly, the running time of this DP algorithm is $\mathcal{O}(T^2)$.

REFERENCES

- [1] N. Absi, S. Dauzère-Pérès, S. Kedad-Sidhoum, B. Penz, and C. Rapine, Lot-sizing with carbon emission constraints, *Eur J Oper Res* 227(1) (2013), 55–61.
- [2] A. Aggarwal and J. K. Park, Improved algorithms for economic lot-size problems, *Oper Res* 14 (1993), 549–571.
- [3] S. Benjaafar, Y. Li, and M. Daskin, Carbon footprint and the management of supply chains: Insights from simple models, *IEEE Trans Autom Sci Eng* 10(1) (2013), 99–116.
- [4] G. R. Bitran and H. H. Yanasse, Computational complexity of the capacitated lot size problem, *Manag Sci* 28 (1982), 1174–1186.
- [5] R. Blanquero and E. Carrizosa, A D.C. biobjective location model, *J Global Optim* 23 (2002), 139–154.
- [6] Y. Bouchery, A. Ghaffari, Z. Jemai, and Y. Dallery, Including sustainability criteria into inventory models, *Eur J Oper Res* 222(2) (2012), 229–240.
- [7] C. S. Chung and C. H. M. Lin, An $\mathcal{O}(T^2)$ algorithm for the NI/G/NI/ND capacitated lot size problem, *Manag Sci* 34 (1988), 420–426.
- [8] M. Ehrgott and X. Gandibleux, A survey and annotated bibliography of multiobjective combinatorial optimization, *OR Spektrum* 22 (2000), 425–460.
- [9] M. Ehrgott and X. Gandibleux, Approximative solution methods for multiobjective combinatorial optimization, *TOP*, 12(1) (2004), 1–89.
- [10] A. Federgruen and J. Meissner, Competition under time-varying demands and dynamic lot-sizing costs, *Nav Res Log* 56(1) (2009), 57–73.

- [11] A. Federgruen and M. Tzur, A simple forward algorithm to solve general dynamic lot sizing models with n periods in $\mathcal{O}(n \log n)$ or $\mathcal{O}(n)$ time, *Manag Sci* 37 (1991), 909–925.
- [12] M. Florian and M. Klein, Deterministic procurement planning with concave costs and capacity constraints, *Manag Sci* 18 (1971), 12–20.
- [13] M. Florian, J.K. Lenstra, and A.H.G. Rinnooy Kan, Deterministic production planning: algorithms and complexity, *Manag Sci* 26 (1980), 669–679.
- [14] M.R. Garey and D.S. Johnson, *Computers and intractability: A guide to the theory of NP-completeness*, W.H. Freeman and Company, New York, 1979.
- [15] S. M. Gilbert, Coordination of pricing and multiple-period production for constant priced goods, *Eur J Oper Res* 114 (1999), 330–337.
- [16] F. Love, Bounded production and inventory models with piecewise concave costs, *Manag Sci* 20 (1973), 313–318.
- [17] S. Mittal and A.S. Schulz, A general framework for designing approximation schemes for combinatorial optimization problems with many objectives combined into one, *Oper Res* 61(2) (2013), 386–397.
- [18] C.H. Papadimitriou and M. Yannakakis, “On the approximability of trade-offs and optimal access of web sources,” In: *Proceedings of 41st Annual Symposium on Foundations of Computer Science*, Los Alamitos, California, 2000, pp. 86–92.
- [19] M.J. Retel Helmrich, R. Jans, W. van den Heuvel, and A.P.M. Wagelmans, The economic lot-sizing problem with an emission constraint, *Forthcoming in Eur J Oper Res*, (2014).
- [20] T. Stidsen, K.A. Andersen, and B. Dammann, A branch and bound algorithm for a class of biobjective mixed integer programs, *Manag Sci* 60(4) (2014), 1009–1032.
- [21] E. Toczyłowski, “An $\mathcal{O}(T^2)$ algorithm for the lot-sizing problem with limited inventory levels,” in: *IEEE Symposium on Emerging Technologies & Factory Automation*, Vol. 3, Los Alamitos, California, 1995, pp. 78–85.
- [22] E.L. Ulungu and J. Teghem, Multi-objective combinatorial optimization problems: a survey, *J Multi-Criteria Decis Anal* 3 (1994), 83–104.
- [23] C.P.M. van Hoesel and A.P.M. Wagelmans, An $\mathcal{O}(T^3)$ algorithm for the economic lot-sizing problem with constant capacities, *Manag Sci* 42 (1996), 142–150.
- [24] C.P.M. van Hoesel and A.P.M. Wagelmans, Parametric analysis of setup cost in the economic lot-sizing model without speculative motives, *Int J Prod Econ* 66 (2000), 13–22.
- [25] J.C. Velázquez-Martínez, J.C. Fransoo, E.E. Blanco, and J. Mora-Vargas, The impact of carbon footprinting aggregation on realizing emission reduction targets, *Flex Serv Manuf J* 26(1–2) (2014), 196–220.
- [26] A.P.M. Wagelmans, C.P.M. van Hoesel, and A. Kolen, Economic lot sizing: An $\mathcal{O}(n \log n)$ algorithm that runs in linear time in the Wagner-Whitin case, *Oper Res* 40 (1992), S145–S156.
- [27] H.M. Wagner and T.M. Whitin, Dynamic version of the economic lot size model, *Manag Sci* 5 (1958), 89–96.
- [28] L.A. Wolsey, Lot-sizing with production and delivery time windows, *Math Program Ser A* 107 (2006), 471–489.
- [29] W.I. Zangwill, Minimum concave cost flows in certain networks, *Manag Sci* 14 (1968), 249–265.