Int. J. Production Economics 145 (2013) 53–66

Contents lists available at SciVerse ScienceDirect

Int. J. Production Economics

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

Dynamic lot sizing for a warm/cold process: Heuristics and insights

Ayhan Özgür Toy [a,*], Emre Berk [b]

[a] Faculty of Industrial Engineering Department, Istanbul Bilgi University, 34060 Eyüp, Istanbul, Turkey
[b] Faculty of Business Administration, Bilkent University, 06800 Bilkent, Ankara, Turkey

ARTICLE INFO

Article history:
Received 15 February 2012
Accepted 10 September 2012
Available online 23 September 2012

Keywords:
Lot sizing
Warm/cold process
Rolling horizon
Heuristics

ABSTRACT

We consider the dynamic lot sizing problem for a warm/cold process where the process can be kept warm at a unit variable cost for the next period if more than a prespecified quantity has been produced. Exploiting the optimal production plan structures, we develop nine rule-based forward solution heuristics. Proposed heuristics are modified counterparts of the heuristics developed previously for the classical dynamic lot sizing problem. In a numerical study, we investigate the performance of the proposed heuristics and identify operating environment characteristics where each particular heuristic is the best or among the best. Moreover, for a warm/cold process setting, our numerical studies indicate that, when used on a rolling horizon basis, a heuristic may also perform better costwise than a solution obtained using a dynamic programming approach.

1. Introduction

In this paper, we consider the problem of dynamic lot sizing for a special type of production processes. The dynamic lot sizing problem is defined as the determination of the production plan which minimizes the total (fixed setup, holding and variable production) costs incurred over the planning horizon for a storable item facing known demands.

Recently, the notion of a "warm/cold process" has been introduced into the scheduling literature (Toy and Berk, 2006). A warm/cold process is defined as a production process that can be kept *warm* for the next period if a minimum amount (the so-called warm threshold) has been produced in the current period and would be *cold*, otherwise. Production environments where the physical nature of the production technology dictates that the processes be literally kept warm in certain periods to avoid expensive shutdown/startups are typical in glass, steel and ceramic production. Robinson and Sahin (2001) provide other examples in food and petrochemical industries where certain cleanup and inspection operations can be avoided in the next period if the quantity produced in the current period exceeds a certain threshold. Production processes where production rates can be varied also fall into the warm/cold process category. The upper bound on the production rate is the physical capacity of the production process and the lower bound corresponds to the warm threshold, below which the process cannot be kept running into the next period without incurring a setup. Such variable production rates can be found in both discrete item manufacturing and process industries. Change in production rate can be obtained at either zero or positive cost depending on the characteristics of the employed technology. The additional variable cost is, then, the variable cost of keeping the process warm onto the next period.

As the above examples illustrate, the dynamic lot-sizing problem in the presence of production quantity—dependent warm/cold processes is a common problem. This problem, in the presence of no shortages, has been formulated and solved optimally by Toy and Berk (2006) using a dynamic programming approach with an $O(N^3)$ forward algorithm where N denotes the problem horizon length. Later, they extend their results to the case where some of the demands may be lost under a profit maximization objective (Berk et al., 2008).

The dynamic lot sizing problem for a warm/cold process is a generalization of the so-called classical problem which was first analyzed by Wagner and Whitin (1958). The classical problem assumes uncapacitated production and no shortages. Wagner and Whitin (1958) provide a dynamic programming solution algorithm and structural results on the optimal solution. Their fundamental contribution lies in establishing the existence of planning horizons, which makes forward solution algorithms possible. Although the optimal solution structure is known, the complexity of obtaining it (shown to be $O(N \log N)$ in general by Federgruen and Tzur, 1991; Wagelmans et al., 1992; Feng et al., 2011 for constant capacities) has stimulated a stream of research that focuses on developing lot sizing heuristics based on simple stopping rules, such as Silver–Meal (Silver and Meal, 1973), Part-Period Balancing (DeMatteis, 1968), Least Unit Cost, Economic Order Interval, McLauren's Order Moment (Vollmann et al., 1997), Least Total Cost (Narasimhan and McLeavy, 1995), Groff's

* Corresponding author.
E-mail addresses: ozgur.toy@bilgi.edu.tr (A.Ö. Toy), eberk@bilkent.edu.tr (E. Berk).

0925-5273/$ - see front matter © 2012 Elsevier B.V. All rights reserved.
http://dx.doi.org/10.1016/j.ijpe.2012.09.006

Algorithm (Groff, 1979). (See also Sahin et al., 2008; Narayanan and Robinson, 2010.)

Further results on the lot sizing problem are found in the literature on its extension to the capacitated production settings. The capacitated lot sizing problem (CLSP) is related to the lot sizing problem for a warm/cold process under certain conditions (see Toy and Berk, 2006). The CLSP has been first studied by Manne (1958) and has been shown to be NP-hard by Florian et al. (1980). Reviews of the works on CLSP (along with the uncapacitated versions) are by Brahimi et al. (2006) and Quadt and Kuhn (2008), who include extensions of the problem, and Buschkühl et al. (2010). Recent analytical studies have focused on novel solution approaches. Heuvel and Wagelmans (2006) develop an $O(T^2)$ algorithm. Pochet and Wolsey (2010) provide a mixed integer programming reformulation that can be solved with LP-relaxation to optimality under reasonable conditions. Chubanov et al. (2008) and Ng et al. (2010) introduce polynomial approximations. Hardin et al. (2007) analyze the quality of bounds by fast algorithms. Reviews of meta-heuristic approaches to the CLSP can be found in Staggemeier and Clark (2001), Jans and Degraeve (2007) and in Guner Goren et al. (2010) on genetic algorithms for lot sizing. A recent review of related works appears also in Glock (2010).

Rule-based heuristics in rolling horizon environments have been studied by Stadtler (2000), Simpson (2001), and Heuvel and Wagelmans (2005). The work herein joins this stream by considering the dynamic lot sizing problem for a warm/cold process. Specifically, we propose rule-based lot sizing heuristics for the problem and examine the efficacy of such rules. To the best of our knowledge, this is the first work that studies lot sizing rules for the operating environment where the production process can be kept warm at some cost if production quantity in a period exceeds a threshold value. We believe that our contributions lie in developing a number of heuristics which perform well in certain operational environments and in identifying such regions for selecting a particular heuristic. We consider the application of the proposed heuristics in a static setting as well as on a rolling horizon basis as it is the practice. The available commercial ERP software (*e.g.*, *SAP*) still offer well-known heuristics for the classical lot sizing problem as options for decision-makers along with the 'optimal' solution algorithms in their manufacturing modules. For the conventional production environments, the benefits of heuristics include the ease of use, smoother production schedules and more intuition for the trade-offs. Moreover, for a warm/cold process setting, our numerical studies indicate that, when used on a rolling horizon basis, a heuristic may also perform better costwise than a solution obtained using a dynamic programming approach. This finding is consistent with similar studies on the classical problem (Stadtler, 2000; Heuvel and Wagelmans, 2005). Hence, investigation of heuristics for warm/cold process settings may be financially beneficial in practice as well as from a purely theoretical perspective. Our work extends the heuristics literature on the dynamic lot sizing problem.

The rest of the paper is organized as follows: In Section 2, we introduce the basic assumptions of our model, formulate the optimization problem and present some key results. In Section 3, we present some theoretical results on an economic production quantity (EPQ) model that we use as a continuous counterpart of a warm/cold process to develop some of our heuristics. In Section 4, we introduce and construct nine lot sizing heuristics for a warm/cold process. In Section 5, we present a numerical study and discuss our findings in regards to the cost performance of the proposed heuristics. In our numerical study, we provide results on the performance distribution of individual heuristics, on the rankings of the heuristics, on identifying the operating environment where a particular heuristic may perform best and on the impact of planning horizon lengths when production plans are made and executed on a rolling horizon basis.

2. Model: assumptions and formulation

We consider the operational setting in Toy and Berk (2006) with time-invariant system and cost parameters. We assume that the length of the problem horizon, N is finite and known. Demand in period t, denoted by D_t $(t=1,2,\dots,N)$, is non-negative and known, but may be different over the problem horizon. No shortages are allowed; that is, the amount demanded in a period has to be produced in or before its period. The amount of production in period t is denoted by x_t. If $x_t>0$, the production indicator δ_t is 1, zero otherwise. The inventory on hand at the end of period t is denoted by y_t $(=y_{t-1}+x_t-D_t)$. Inventory holding cost per unit of ending inventory is h per period. Without loss of generality, we assume that the initial inventory level is zero. We assume that unit production cost is c but may be omitted in the analysis since all demands must be met over the horizon.

Production quantity in a period cannot exceed the capacity, R. For feasibility, we assume that, for any t, there exists a $j(t)$ for which $\sum_{i=t}^{j(t)} D_i \le (j(t)-t+1)R$ for $t \le j(t) \le N$, $1 \le t \le N$. This condition guarantees that any subset of demands can be produced within the horizon; a special case of the condition is satisfied when $D_t \le R$ for all t. We consider a warm/cold production process: The production process may be kept *warm* onto the beginning of period t if the production quantity in the previous period is at or above a threshold value Q; that is, $x_{t-1} \ge Q$. Otherwise, the process cannot be kept warm and is *cold*. Let z_t indicate the warm/cold status of the process as period t starts; it attains a value of 0 if the process is warm and 1, otherwise. In order to keep the process warm onto period t, warming cost ω is charged for every unit of unused capacity in period $t-1$. That is, the warming cost incurred in period $t-1$ would be $\omega(R-x_{t-1})$ monetary units. Note that, even if the quantity produced in period $t-1$ is at least Q, it may not be optimal to keep the process warm onto the next period if there would not be any production during the next period. In such instances, there will be no warming costs incurred although $x_{t-1} \ge Q$ since $x_t=0$. We assume that a warm process requires no setup (and, hence incurs no setup cost) but a cold process requires a cold setup with a fixed cost K (>0) if production is to be done in the period. Finally, we assume that $h>\omega$ which ensures the Wagner–Whitin type cost structure, and that the warm/cold process threshold is between the point of indifference and the capacity, $R-(K/\omega)<Q\le R$. (For the implications of these assumptions, see Toy and Berk, 2006.)

The objective is to find a production plan $x_t \ge 0$ $(t=1,2,\dots,N)$ (timing and amount of production), such that all demands are met at minimum total cost over the horizon. Let $X=\{x_1,\dots,x_N\}$ denote a feasible production plan constructed over periods 1 through N; Γ_t be the variable cost incurred within period t computed as $\Gamma_t = hy_t+\omega(R-x_t)\delta_{t+1}(1-z_{t+1})$ under the given production plan; and THC denote the total horizon cost. Then, the optimization problem (P) can be formally stated as follows:

$$\min_X THC = \sum_{t=1}^{N} (K\delta_t z_t + \Gamma_t)$$

subject to $0 \le x_t \le R$, $y_{t-1}+x_t \ge D_t$, for all t; $z_1=1$, $\delta_{N+1}=0$, $z_{N+1}=1$, and $y_0=y_N=0$.

Let $L_{u,v}$ represent a subset of X between periods u and $v-1$ (inclusive) such that the starting inventory in period u and ending inventory in period $v-1$ are zero and production is done in all periods u through m to cover the demands for periods u through $v-1$. Formally, $L_{u,v}=\{x_t \mid x_t>0, t=u,\dots,m;\ x_t=0 \text{ for } t=m+1,\dots,v-1;\ y_{u-1}=y_{v-1}=0\}$ for $0 \le u \le m < v \le N+1$. With a slight abuse of terminology, we shall refer to $L_{u,\nu}$ as a *production lot* (for period u through $\nu-1$), the total quantity produced in the production lot as its *lot size* and $(\nu-u)$ as the *production lot run length*. Under the assumed cost structure, we have the following result:

Proposition 1 (*Optimal Schedule within a Production Lot*). *The optimal* (*total cost minimizing*) *structure of a production lot $L_{u,\nu}$ is as follows*: (i) $x_t = \max(Q,[D_t - y_{t-1}])$, *for* $u \le t \le m-1$ *and* $m-u > 0$, (ii) $x_m = \varepsilon$, *where* $\varepsilon = [\sum_{i=m}^{\nu-1} D_i - y_{m-1}]^+ < R$ *and* $m-u \ge 0$, *and* (iii) $x_t = 0$ *for* $t = m+1, \ldots, \nu-1$.

The proof rests on showing that any other production plan for the periods u through $\nu-1$ with the given zero inventory constraints would result in higher total costs and hence cannot be optimal. Note that, typically, in the uncapacitated, classical setting, each production lot corresponds to a single period of production followed by non-production periods (with $u=m$); only in rare cases where setup costs are (comparably) very small, one would get a lot-for-lot schedule. But, for a warm/cold process, a production may contain a number of periods in which production is done in succession (with $u \le m$). Also note that the optimal production plan for problem (*P*) consists of production lots that have the same structure as that given above (see Theorems 1–4 in Toy and Berk, 2006).

Similarly, in devising and implementing the forward solution heuristics below, we will use this key result. All production lots $L_{u,\nu}$ will be assumed to have the above optimal structure. Hence, heuristics will result in lot sizes (*i.e.*, the selection of some period u and ν) that may be sub-optimal but the structure of the production plan will conform to the optimal solution structure.

Let $C_{u,\nu}$ be the variable cost incurred over periods u through $\nu-1$ by the production lot $L_{u,\nu}$,

$$C_{u,\nu} = \sum_{t=u}^{\nu-1} \Gamma_t = \sum_{t=u}^{\nu-1} h y_t + \sum_{t=u}^{m-1} \omega(R-x_t)$$

In the following Lemma, we state how incremental updating of production plans and costs is performed as the production lot length is extended by one period, *i.e.*, how the production lot $L_{u,\nu+1} = \{x'_t \,|\, x'_t > 0,\ t=u,\ldots,m';\ x'_t = 0 \text{ for } t=m'+1,\ldots,\nu-1\}$ and the corresponding cost $C_{u,\nu+1}$ are obtained dynamically from $L_{u,\nu} = \{x_t \,|\, x_t > 0, t=u,\ldots,m;\ x_t = 0 \text{ for } t=m+1,\ldots,\nu-1\}$ and $C_{u,\nu}$, respectively.

Lemma 1. *Let* $L_{u,\nu} = \{x_t \,|\, x_t > 0, t=u,\ldots,m;\ x_t = 0 \text{ for } t=m+1,\ldots,\nu-1\}$, $y'_t = y'_{t-1} + x'_t - D_t$ *and* $E_t = (\max(Q,[D_t - y'_{t-1}]^+)(h-\omega) + R\omega)/h$.

(a) $L_{u,\nu+1} = \{x'_t \,|\, x'_t > 0, t=u,\ldots,m';\ x'_t = 0 \text{ for } t=m'+1,\ldots,\nu-1\}$ *where*

$$x'_t = \begin{cases} x_t & \text{for } u \le t < m \\ D_\nu + x_m & \text{for } t=m \text{ if } D_\nu + x_m \le E_t \\ \max(Q,[D_t - y'_{t-1}]^+) & \text{for } t=m \text{ if } D_\nu + x_m > E_t \\ \left[D_\nu + x_m - \sum_{j=m}^{t-1} x'_j\right]^+ & \text{for } m < t < \nu+1 \\ & \text{if } \left[D_\nu + x_m - \sum_{j=m}^{t-1} x'_j\right]^+ \le E_t \\ \max(Q,[D_t - y'_{t-1}]^+) & \text{for } m < t < \nu+1 \\ & \text{if } \left[D_\nu + x_m - \sum_{j=m}^{t-1} x'_j\right]^+ > E_t \end{cases}$$

and $m' = \max\{t : x'_t > 0, u \le t < \nu+1\}$;

(b) $C_{u,\nu+1} = C_{u,\nu} + \sum_{i=m}^{\nu-1} h(\sum_{j=m}^{i} x'_j - x_m) + \sum_{i=m}^{m\prime-1} \omega(R - x'_i)$.

Proof. (a) Immediately follows from the optimal schedule structure.

(b) Let $C_{u,\nu}$ be the total variable cost associated with the production lot $L_{u\nu}$ as defined before.

Hence,

$$\begin{aligned} C_{u,\nu+1} - C_{u,\nu} &= \left[\sum_{i=u}^{\nu} h \sum_{j=u}^{i} (x'_j - D_j) + \sum_{i=u}^{\nu} I_{(x'_i>0)(x'_{i+1}>0)} \omega(R - x'_i)\right] \\ &\quad - \left[\sum_{i=u}^{\nu-1} h \sum_{j=u}^{i} (x_j - D_j) + \sum_{i=u}^{m-1} \omega(R - x_i)\right] \\ &= \left[\sum_{i=u}^{m-1} h \sum_{j=u}^{i} (x'_j - D_j) + h \sum_{j=u}^{m} (x'_j - D_j)\right. \\ &\quad + \sum_{i=m+1}^{\nu-1} h \sum_{j=u}^{i} (x'_j - D_j) + h \sum_{j=u}^{\nu} (x'_j - D_j) \\ &\quad \left. + \sum_{i=u}^{m-1} \omega(R - x'_i) + \sum_{i=m}^{\nu} I_{(x'_i>0)(x'_{i+1}>0)} \omega(R - x'_i)\right] \\ &\quad - \left[\sum_{i=u}^{m-1} h \sum_{j=u}^{i} (x_j - D_j) + h \sum_{j=u}^{m} (x_j - D_j)\right. \\ &\quad \left. + \sum_{i=m+1}^{\nu-1} h \sum_{j=u}^{i} (x_j - D_j) + \sum_{i=u}^{m-1} \omega(R - x_i)\right] \\ &= \sum_{i=m}^{\nu} h \left(\sum_{j=m}^{i} x'_j - x_m\right) + \sum_{i=m}^{\nu} I_{(x'_i>0)(x'_{i+1}>0)} \omega(R - x'_i) \end{aligned}$$

which reduces to the result. □

Before introducing the heuristics, we will explore some properties of the continuous review economic production quantity (EPQ) counterpart of a warm/cold process in the next section.

3. Economic production quantity for a warm/cold process

In this section, we construct the continuous review economic production quantity (EPQ) counterpart of a warm/cold process. The model derived herein serves as the foundation of the EPQ-based heuristics to be discussed later. The model assumptions are as follows: Demand rate is deterministic and constant, d. There is a constant production rate, p which is a decision variable and may take on values over $[\max(Q,d),R]$ per unit time where Q denotes the physical threshold for a warm process. The cost rate of operating with a particular value of the production rate p_0 is given by $\omega(R-p_0)$. This cost is analogous to the warming cost in the sense that it keeps the production process running. We omit the unit production cost due to material usage, *etc.* as in the periodic review problem at hand. Hence, the effective production cost rate is $\omega(R-p_0)$. A cycle is defined as the time between two consecutive instances of a cold setup, which initiates a production run. Each cycle consists of two phases – a production and a non-production phase – as in the classical EPQ setting. During the production phase, q units are produced and inventory is accumulated at a rate that is in excess of the demand rate; during the non-production phase, the accumulated inventory is used up to satisfy the demand until the next production run. Each unit of inventory held per unit time incurs a carrying cost h. The objective is to find the production quantity and the production rate that minimize the total cost per unit time.

The cycle length for a given production quantity q and a production rate p is given by

$$CL(q,p) = q/d$$

The length of the production phase in a cycle is q/p and the maximum inventory level at the end of the production phase is $(p-d)(q/p)$. In each cycle, there is a single cold setup by definition, and warming cost is incurred during the entire production phase. Therefore, the total (setup, holding and production/warming) cost per cycle is

$$CC(q,p) = K + \frac{q^2(p-d)}{2pd}h + \frac{q(R-p)}{p}\omega$$

The total cost rate is $TC(q,p) = CC(q,p)/CL(q,p) = Kd/q + (q(p-d)/2p)h + [\omega(R-p)/p]d$. Note that the cost rate expression is similar in structure to that of an economic production quantity model in the presence of (i) production rate dependent unit production costs *or* (ii) cost of selecting a production rate (e.g. Khouja and Mehrez, 1994; Larsen, 1997).The unique optimizer $\hat{q}(p)$ of $TC(q,p)$ for a given value of production rate p is given by the classical EPQ formula: $\hat{q}(p) = \sqrt{2Kd/h}\sqrt{p/(p-d)}$ with the corresponding optimal cost rate $\widehat{TC}(p) = \sqrt{2Kdh}\sqrt{(p-d)/p} + d\omega(R-p)/p$. (Note, when $p=d$, the process is run continuously ($\hat{q}(p)\to\infty$) resulting in $\widehat{TC}(p=d) = \omega(R-d)$.) The unique extremum of $\widehat{TC}(p)$ is given by $\tilde{p} = 2\omega^2R^2d/(2\omega^2R^2 - dKh)$. However, it is not possible to say that $\tilde{p}$ is a global minimum (or maximum) since the principal minors of the Hessian are of mixed signs. Furthermore, it may not lie within the feasible region. Therefore, to obtain the optimal production rate p^*, we need to check the total cost rates attained when the production rate takes on the extremum value (if feasible) and its two boundaries; $p^* = \arg\min_{p\in\{\max(Q,d),R\}\cup\hat{P}} \widehat{TC}(p)$ where we define the set $\hat{P} = \{r : r = \tilde{p}; \max(Q,d) \le r \le R\}$ which is non-empty only for feasible $\tilde{p}$. Then, we have $\widehat{TC^*} = \widehat{TC}(p^*)$ and $\widehat{q^*} = \hat{q}(p^*)$; we also let $\widehat{CC^*}$ denote the corresponding cycle cost.

4. Heuristics

In this section, we develop the proposed lot sizing heuristics. Solving the problem (*P*) over a problem horizon from period i' to period N' by a forward heuristic means that a production plan is obtained in a forward manner (starting from period i' and proceeding ahead up to period N') by employing a pre-specified stopping rule to determine the production lots over the horizon. The stopping rule of a heuristic dictates when a production lot starting in period u should terminate; at the termination of one production lot, it is assumed that another starts. Due to Proposition 1, the structure of the production lot is known. Applying the rule over the horizon to generate successive production lots, one obtains the production plan X under the heuristic. The total horizon cost *THC* under the heuristic is then computed using the production plan X. The implementation of the forward heuristics herein follows the pseudo-code below.

```
Program_FindProductionPlan(i',N')
begin
    i := i'
    while (i ≤ N')
        u := i;
        n := 1;
        stoping_rule :=false
        generate L_{u,u+n}
        while (stopping_rule==false)
            generate L_{u,u+n+1}
            if stopping condition is satisfied then do
                stopping_rule :=true
                compute n̂ (number of periods in the
                production lot)
                i := u+n̂
            else n := n+1
end
```

We propose nine heuristics for the dynamic lot sizing problem for a warm/cold process. They are adaptations or modifications for the warm/cold process setting of the heuristics that have been developed for the classical, uncapacitated lot sizing problem. As we show in our numerical analysis, some of the proposed heuristics do not perform as well for certain system and cost parameters and demand patterns. However, there is not one particular heuristic that performs best in all experiments either. Below, we discuss the construction of our heuristics and state the stopping rules for each heuristic.

Heuristic #1 is developed in the same essence as the Silver–Meal heuristic in the classical, uncapacitated setting (Silver and Meal, 1973); it is an adaptation of the Silver–Meal heuristic for a warm/cold process. It rests on the comparison of the total cost per period. A production lot starting at u terminates at period $u+\hat{n}-1$ after which the total cost rate increases for the first time. Formally, the stopping rule for this heuristic can be stated as $K + C_{u,u+n}/n < K + C_{u,u+n+1}/(n+1)$ with $\hat{n} = 1 + \max\{n : K + C_{u,u+n}/n \ge K + C_{u,u+n+1}/(n+1)\}$.

To illustrate the mechanics of the proposed heuristics, we shall use an illustrative example with the following parameter set: $K=10$, $h=1$, $\omega = 0.95$, $R=8$, and $Q=5$ with $N=15$ and demands over the horizon, $\mathbf{D} = \{6,4,2,2,4,7,5,6,4,8,4,5,8,1,4\}$. The optimal production plan for periods 1 through 15 for this example is found as {{6,5,3,0}, {5,6,5,6,5,7,5,5,7,5,0}} with a corresponding total cost of 55.70. Note that the production plan is presented as a sequence of production lots and each production lot is defined as a sequence of the production done on each consecutive period. For Heuristic #1, starting with the first period, we construct the production lot $L_{u,u+n}$, compute the stopping criterion $\rho_{u,u+n} = (K + C_{u,u+n})/n$ and apply the rule as follows. $L_{12} = \{6\}$, $\rho_{12} = (10+0)/(2-1) = 10$; $L_{13} = \{6,4\}$, $\rho_{13} = (10+0.95(8-6))/(3-1) = 5.95$ $(<\rho_{12})$; $L_{14} = \{6,6,0\}$, $\rho_{14} = (10+0.95(8-6)+1(0+6-4))/(4-1) = 4.63$ $(<\rho_{13})$; $L_{15} = \{6,5,3,0\}$, $\rho_{15} = (10+[0.95(8-6)+0.95(8-5)]+[1(0+5-4)+1(1+3-2)])/(5-1) = 4.44(<\rho_{14})$; $L_{16} = \{6,5,7,0,0\}$, $\rho_{16} = 5.15$ $(>\rho_{15})$. We freeze the production plan for periods 1 through 4, {6,5,3,0}, and start from period 5. Proceeding in a similar fashion, we get the production plan for periods 1 through 15 as {{6,5,3,0}, {5,6,5,6,5,7,4}, {5,8,5,0}} with a corresponding total cost of 59.90.

Heuristic #2 is similar in construction to the Part Period Balancing (PPB) heuristic in the classical, uncapacitated setting (DeMatteis, 1968); it is an adaptation of the PPB heuristic for a warm/cold process. It rests on the comparison of the cost of a production lot against a cold setup. A production lot starting at u terminates at period $u+\hat{n}-1$ after which the variable cost exceeds the fixed cold setup cost for the first time. Formally, the stopping rule can be stated as $C_{u,u+n} > K$ with $\hat{n} = \max\{n : C_{u,u+n} \le K\}$.

Consider the same illustrative setting given above. Starting with the first period, we construct the production lot $L_{u,u+n}$, compute the stopping criterion $\rho_{un} = C_{u,u+n}$ and apply the rule as follows. $L_{12} = \{6\}$, $\rho_{12} = 0$; $L_{13} = \{6,4\}$, $\rho_{13} = 0.95(8-6) = 1.90$ $(<K=10)$; $L_{14} = \{6,6,0\}$, $\rho_{14} = 0.95(8-6) + 1(0+6-4) = 3.90$ $(<K)$; $L_{15} = \{6,5,3,0\}$, $\rho_{15} = [0.95(8-6)+0.95(8-5)] + [1(0+5-4)+1(1+3-2)] = 7.75(<K)$; $L_{16} = \{6,5,7,0,0\}$, $\rho_{16} = 15.75$ $(>K)$. We freeze the production plan for periods 1 through 4, {6,5,3,0}, and start from period 5. Proceeding in a similar fashion, we get the production plan for periods 1 through 15 as {{6,5,3,0}, {5,6,5,6}, {5,7,5,4}, {8,5,0}} with a corresponding total cost of 60.90.

Heuristic #3 is the adaptation for a warm/cold process of the Least Unit Cost heuristic (LUC) in the classical, uncapacitated setting (Vollmann et al., 1997). It rests on the total cost per unit comparison. A production lot starting at u terminates at period $u+\hat{n}-1$ after which the total cost rate increases for the first time. Formally, the stopping rule can be stated as $(K+C_{u,u+n})/\sum_{i=u}^{u+n-1} D_i < (K+C_{u,u+n+1})/\sum_{i=u}^{u+n} D_i$ with $\hat{n} = 1+\max\{n : (K+C_{u,u+n})/\sum_{i=u}^{u+n-1} D_i \geq (K+C_{u,u+n+1})/\sum_{i=u}^{u+n} D_i\}$.

For Heuristic #3, starting with the first period, we construct the production lot $L_{u,u+n}$, compute the stopping criterion $\rho_{un} = (K+C_{u,u+n})/\sum_{i=u}^{u+n-1} D_i$ and apply the rule as follows. $L_{12} = \{6\}$, $\rho_{12} = (10+0)/(6) = 1.67$; $L_{13} = \{6,4\}$, $\rho_{13} = (10+0.95(8-6))/(6+4) = 1.19$ $(<\rho_{12})$; $L_{14} = \{6,6,0\}$, $\rho_{14} = (10+0.95(8-6)+1(0+6-4))/(6+4+2) = 1.16$ $(<\rho_{13})$; $L_{15} = \{6,5,3,0\}$, $\rho_{15} = (10+[0.95(8-6)+0.95(8-5)]+[1(0+5-4)+1(1+3-2)])/(6+4+2+2) = 1.27$ $(>\rho_{14})$. We freeze the production plan for periods 1 through 3, {6,6,0}, and start from period 4. Proceeding in a similar fashion, we get the production plan for periods 1 through 15 as {{6,6,0}, {5,5,5,5,5,5,6,5,5,7,1},{4}} with a corresponding total cost of 75.55.

Heuristic #4 is the adaptation for a warm/cold process of the Least Total Cost (LTC) heuristic in the classical, uncapacitated setting (Narasimhan and McLeavy, 1995). A production lot starting at u terminates at period $u+\hat{n}-1$ at which the absolute difference between the fixed cold setup cost and the total variable cost is minimum. Formally, the stopping rule can be stated as $C_{u,u+n} > K$ with $\hat{n} = \arg\min_{n \in \{\tilde{n},\tilde{n}+1\}} |K - C_{u,u+n}|$ where $\tilde{n} = \max\{i : C_{u,u+i} \leq K\}$.

Starting with the first period, we construct the production lot $L_{u,u+n}$, compute the stopping criterion $\rho_{un} = C_{u,u+n}$ and apply the rule as follows. $L_{12} = \{6\}$, $\rho_{12} = 0$; $L_{13} = \{6,4\}$, $\rho_{13} = 0.95\ (8-6) = 1.90$ $(<K=10)$; $L_{14} = \{6,6,0\}$, $\rho_{14} = [0.95(8-6)+1(0+6-4)] = 3.90$ (<10); $L_{15} = \{6,5,3,0\}$, $\rho_{15} = [0.95(8-6)+0.95(8-5)]+[1(0+5-4)+1(1+3-2)] = 7.80$ (<10); $L_{16} = \{6,5,7,0,0\}$, $\rho_{16} = 15.75$ (>10). Note that the difference for period 5 is larger than that for period 4; $\min(|10-15.75|,|10-7.80|) = 2.20$. Hence, we freeze the production plan for periods 1 through 4, {6,5,3,0}, and start from period 5. Proceeding in a similar fashion, we get the production plan for periods 1 through 15 as {{6,5,3,0}, {5,6,5,6,4},{8,5,5,7,1},{4}} with a corresponding total cost of 66.95.

For the classical, uncapacitated setting, a group of heuristics have been proposed in the literature based on a continuous review (economic order quantity, EOQ) approximation of the periodic review production environment: Groff's heuristic (Groff, 1979), the economic order interval (EOI) heuristic, and McLauren's Order Moment (MOM) heuristic (see Vollmann et al., 1997). In these heuristics, the demand stream over the problem horizon is approximated by a constant rate demand process with the same mean; and, the cost trade-offs inherent in the lot sizing problem for the periodic setting are approximated by the cost trade-offs based on the EOQ model. This class of heuristics has been shown to be effective in certain operating environments. For the warm/cold production processes, we propose four heuristics that are adaptations or modifications of these classical setting heuristics. The existence of a warm/cold process implies the capability of keeping the process warm (running) by selecting minimum production quantities (at or above the warm threshold) and incurring a warming cost. This suggests an analogy between the warm/cold setting and a production environment with a finite production rate. Hence, the construction of our heuristics is based on the approximation of the periodic review, warm/cold production setting via the continuous review, economic production quantity (EPQ) model. It attempts to capture the warm/cold capability through the finite production rate that can be selected at a cost. For these heuristics, we have the following definitions. Let $\overline{D}$ denote the constant demand rate for the problem horizon; $\widehat{TC^*}$ denote the optimal total cost rate and $\widehat{q^*}$ denote the optimal economic production quantity for the *equivalent* EPQ model. The equivalent EPQ model assumes constant demand rate $\overline{D}$ and a constant production rate p^* whose computation involves the warm threshold Q, the physical production capacity R and the warming cost ω. We refer the reader to Section 3 for the development of the equivalent EPQ approximation and the derivation of the basic operational entities above ($\widehat{TC^*}$, $\widehat{q^*}$ and p^*) used in the construction of the heuristics. We introduce these four heuristics (*Heuristics* #5–8) below.

Heuristic #5 is the adaptation for a warm/cold process of the Economic Order Interval (EOI) heuristic in the classical, uncapacitated setting (Vollmann et al., 1997). It is based on an equivalent (continuous review) economic production quantity $\widehat{q^*}$ obtained by assuming a constant demand rate $\overline{D}$ over the horizon and imposing a constant production rate p^*. By construction, with this rule, a production plan consists of production lots of equal size (except possibly the last lot). In this heuristic, a production lot starting at u terminates at period $u+\hat{n}-1$ when the production lot reaches a length equal to the largest integer that is smaller than the economic production interval (EPI). The stopping rule can be stated as $n > \widehat{q^*}/\overline{D}$ with $\hat{n} = \lfloor \widehat{q^*}/\overline{D} \rfloor$ with $\widehat{q^*} = N\overline{D}$ if $p^* = \overline{D}$.

For the given demand stream and cost parameters of the illustrative example, we have $\overline{D} = 4.67$, $\tilde{p} = 7.83$, $\widehat{TC}(Q) = 5.15$; $\widehat{TC}(\tilde{p}) = 6.237$; $\widehat{TC}(R) = 6.236$; $\widehat{TC}^* = 5.15$; $p^* = 5.00$ and $\widehat{q^*} = 37.42$ for the heuristics based on the EPQ model. For Heuristic #5, $\hat{n} = \lfloor 37.42/4.67 \rfloor = 8$ which gives the production plan for periods 1 through 15 as {{6,5,5,5,5,5,5,0}, {5,7,5,5,7,5,0}} with a corresponding total cost of 85.60.

In the classical, uncapacitated setting, Groff (1979) proposes a heuristic that uses an EOQ approximation to the dynamic lot sizing problem. It is based on the comparison of marginal changes in the fixed and variable cost components when the length of a production lot is extended by one period. We develop *Heuristic* #6 in a similar fashion using, instead, an EPQ approximation for a warm/cold process. Consider the production lot $L_{u,u+n}$. Let $\tilde{C}_{u,u+n}$ denote the total holding and warming costs within the production lot as approximated by its EPQ counterpart for a constant demand rate $\overline{D}$ over the horizon and imposing a constant production rate p^*; and, let $[K+\tilde{C}_{u,u+n}](N/n)$ denote the corresponding approximate total cost over the horizon as in Groff. We obtain $\tilde{C}_{u,u+n}$ as follows. With constant demand rate over the horizon, $\overline{D}$ and the imposed production rate p^*, the maximum inventory level within the production lot is $I_{\max} = \sum_{i=u}^{u+n-1} D_i(p^*-\overline{D})/p^*$. The cost of the inventory carried within the lot is approximated as $hnI_{\max}/2$. Noting that $\sum_{i=u}^{u+n-1} D_i/p^*$ gives the length of the production run within the lot, the corresponding warming cost within the lot is computed as $(R-p^*)\omega \sum_{i=u}^{u+n-1} D_i/p^*$, which we approximate as $(R-p^*)\omega n\overline{D}/p^*$ by replacing the total demand by its average. Then,

$$\tilde{C}_{u,u+n} = hn \sum_{i=u}^{u+n-1} D_i(p^*-\overline{D})/(2p^*) + (R-p^*)\omega n\overline{D}/p^*$$

Under *Heuristic* #6, a production lot starting at u terminates at period $u+\hat{n}-1$ after which the approximate total cost over the horizon increases for the first time. Formally, the stopping rule can be stated as $[K+\tilde{C}_{u,u+n}](N/n) < [K+\tilde{C}_{u,u+n+1}](N/(n+1))$ which simplifies to $n(n+1)D_{u+n} > [\widehat{q^*}]^2/\overline{D}$ $(=2Kp^*/h(p^*-\overline{D}))$ with $\hat{n} = 1+\max\{n : n(n+1)D_{u+n} < [\widehat{q^*}]^2/\overline{D}\}$ with $\widehat{q^*} = N\overline{D}$ if $p^* = \overline{D}$.

Starting with the first period, we construct the production lot $L_{u,u+n}$, compute the stopping criterion $\rho_{un} = [K+\tilde{C}_{u,u+n}](N/n)$ and apply the rule as follows. $L_{12} = \{6\}$, $\rho_{12} = 192.9$; $L_{13} = \{6,4\}$, $\rho_{13} = 119.9\ (<\rho_{12})$; $L_{14} = \{6,6,0\}$, $\rho_{14} = 95.9\ (<\rho_{13})$; $L_{15} = \{6,5,3,0\}$, $\rho_{15} = 84.4\ (<\rho_{14})$; $L_{16} = \{6,5,7,0,0\}$, $\rho_{16} = 78.9\ (<\rho_{15})$; $L_{17} = \{6,5,5,5,4,0\}$, $\rho_{17} = 77.4\ (<\rho_{16})$; $L_{18} = \{6,5,5,5,5,4,0\}$, $\rho_{18} = 76.33$ $(<\rho_{17})$; $L_{19} = \{6,5,5,5,5,5,4,0\}$, $\rho_{19} = 76.65$ $(>\rho_{18})$. We freeze the production plan for periods 1 through 8, {6,5,5,5,5,4,0}, and start from period 9. Proceeding in a similar fashion, we get the production plan for periods 1 through 15 as {{6,5,5,5,5,4,0}, {6,5,7,5,5,7,5,0}} with a corresponding total cost of 77.65.

Next, we develop the two heuristics (*Heuristics* #7 and #8) that are modifications of the McLauren's Order Moment (MOM) heuristic in the classical, uncapacitated setting (Vollmann et al., 1997). The construction logic of the MOM heuristic is similar to that of the PPB heuristic; it determines the lot size for individual orders by matching the number of accumulated part periods to the number that would be incurred if an order of size equal to EOQ were placed under conditions of constant demand. The construction of *Heuristics* #7 *and* #8 is similar; an EPQ equivalent is used with constant demand rate over the horizon, $\overline{D}$ and the imposed production rate p^* instead of the EOQ. In *Heuristic* #7, we use the total cost within the lot as the basis for the stopping rule. A production lot starting at u terminates at period $u+\hat{n}-1$ after which the total cost within the lot exceeds the benchmark value (cost for an EPQ cycle). It can be formally stated as $K+C_{u,u+n} > \widehat{CC^*}$ with $\hat{n} = \max\{n : K+ C_{u,u+n} \leq \widehat{CC^*}\}$. In *Heuristic* #8, we use the production quantity within the lot as the basis for the stopping rule. A production lot starting at u terminates at period $u+\hat{n}-1$ after which the total demand to be satisfied by the lot exceeds the benchmark value (an EPQ). It can be formally stated as $\sum_{i=u}^{u+n-1} D_i > \widehat{q^*}$ with $\hat{n} = \max\{n : \sum_{i=u}^{u+n-1} D_i \leq \widehat{q^*}\}$ with $\widehat{q^*} = N\overline{D}$ if $p^* = \overline{D}$.

Consider the illustrative example. For Heuristic #7, starting with the first period, we construct the production lot $L_{u,u+n}$, compute the stopping criterion $\rho_{un} = K+C_{u,u+n}$ and apply the rule as follows. $L_{12}=\{6\}$, $\rho_{12}=10$; $L_{13}=\{6,4\}$, $\rho_{13}=10+0.95(8-6)=11.90$ $(<(\widehat{q^*}/\overline{D})\widehat{TC^*}=((37.41)(5.15)/4.67)=41.33)$; $L_{14}=\{6,6,0\}$, $\rho_{14}=10+0.95(8-6)+1(0+6-4)=13.90$ (<41.33); $L_{15}=\{6,5,3,0\}$, $\rho_{15}=10+[0.95(8-6)+0.95(8-5)]+[1(0+5-4)+1(1+3-2)]=17.75$ $(<41,33)$; $L_{16}=\{6,5,7,0,0\}$, $\rho_{16}=25.75$ (<41.33); $L_{17}=\{6,5,5,5,4,0\}$, $\rho_{17}=39.45$ $(<41,33)$; $L_{18}=\{6,5,5,5,5,4,0\}$, $\rho_{18}=48.3$ $(>41,33)$. We freeze the production plan for periods 1 through 7, {6,5,5,5,4,0}, and start from period 8. Proceeding in a similar fashion, we get the production plan for periods 1 through 15 as {{6,5,5,5,4,0}, {5,6,5,7,5,5,7,5,0}} with a corresponding total cost of 71.65. For Heuristic #8, starting with the first period, we construct the production lot $L_{u,u+n}$, compute the stopping criterion $\rho_{un} = \sum_{i=u}^{u+n-1} D_i$ and apply the rule as follows. $L_{12}=\{6\}$, $\rho_{12}=6$ $(<\widehat{q^*}=37.42)$; $L_{13}=\{6,4\}$, $\rho_{13}=6+4=10$ (<37.42); $L_{14}=\{6,6,0\}$, $\rho_{14}=12$ (<14.97); $L_{15}=\{6,5,3,0\}$, $\rho_{15}=14$ (<37.42); $L_{16}=\{6,5,7,0,0\}$, $\rho_{16}=18$ (<37.42); $L_{17}=\{6,5,5,5,4,0\}$, $\rho_{17}=25$ (<37.42); $L_{18}=\{6,5,5,5,5,4,0\}$, $\rho_{18}=30$ (<37.42); $L_{19}=\{6,5,5,5,5,5,5,0\}$, $\rho_{19}=36$ (<37.42); $L_{1,10}=\{6,5,5,5,5,5,5,4,0\}$, $\rho_{1,10}=40$ (>37.42). We freeze the production plan for periods 1 through 9, {6,5,5,5, 5,5,5,0}, and start from period 10. Proceeding in a similar fashion, we get the production plan for periods 1 through 15 as {{6,5,5,5,5,5,5,0}, {5,7,5,5,7,5,0}} with a corresponding total cost of 85.60.

Finally, we propose *Heuristic* #9 which is a modification for a warm/cold process of the Wagner–Whitin solution algorithm in the classical, uncapacitated setting (Wagner and Whitin, 1958). The construction of the heuristic is as follows. Given a warm/cold process with its system and cost parameters and a demand stream; we consider another production process with the same parameters and demands except that it has infinite capacity per period and cannot be kept warm. That is, the two processes are identical except that $R\to\infty$ and $Q\to\infty$ for the latter which corresponds to the classical, uncapacitated setting. Let $\check{X} = \{\check{x}_1, \check{x}_2, \ldots, \check{x}_N\}$ denote the production plan obtained by the Wagner–Whitin algorithm for this uncapacitated production process for the given demands. Under *Heuristic* #9, the production plan $X = \{x_1, x_2, \ldots, x_N\}$ for the warm/cold process at hand is obtained by spreading out the production quantities in $\check{X}$ over the periods so as to incur the lowest possible costs by taking into account the capability of keeping the process warm and the physical capacity level. With this heuristic, a production lot starting at u terminates at period $u+\hat{n}-1$ after which a positive production quantity is encountered for the first time in the Wagner–Whitin solution. The stopping rule can be formally stated as $\check{x}_{u+n} > 0$ with $\check{n} = \min\{n : \check{x}_{u+n} > 0, n > 1\}-1$.

For the illustrative example, we have the Wagner–Whitin solution given by {{12,0,0},{6,0},{12,0},{10,0},{17,0,0},{13,0,0}}. Then, the corresponding production plan for the warm/cold process for periods 1 through 15 is {{6,6,0},{6,0},{7,5}, {6,4}, {8,5,4},{8,5,0}} with a corresponding total cost of 71.45.

Note that costs are computed by considering warming effects for consecutive lots even when the stopping rules generate separate lots (beginning with cold setups); that is, heuristics are used to obtain only the production plan but not the costs. For example, under *Heuristic* #4, the second and third lots generated result in a production run that goes on for five consecutive periods (without a cold setup in between) with the process being kept warm between periods seven and eight.

For ease of reference, we provide a list of the basis of construction, the stopping rule and lot run length for all of our heuristics in Table 1.

5. Numerical study and discussion

In this section, we present and discuss our findings in a numerical study.

We conducted our numerical study to investigate four aspects: (i) percentage deviation from the optimal cost for each heuristic; (ii) dominance of heuristics among themselves; (iii) impact of the parameter values and demand patterns on performances of the heuristics, and finally (iv) impact of planning horizon lengths when production plans are made and executed dynamically on a rolling horizon basis.

For our numerical study, we considered a problem horizon of $N=300$ periods. Demands are generated randomly from a normal distribution with mean $\mu(=500)$ and standard deviation σ with $\sigma/\mu \in \{0, 0.2, 0.5, 1\}$. Demand streams generated have been pre-processed: (i) All demand values have been truncated to integers, and negative demands have been replaced by zero demands. (ii) Given the demand stream from (i) $D' = \{D'_1, \ldots, D'_N\}$, the final demand stream $D = \{D_1, \ldots, D_N\}$ is obtained by spreading out D' over the problem horizon starting with the last period such that no final demand value exceeds the imposed capacity for a particular experiment; that is, $D_N = \min(R, D'_N)$ and $D_{N-j} = \min(R, (D'_{N-j} + [\sum_{i=N-j+1}^{N} (D'_i - D_i)]))$ for $1 \leq j \leq N-1$. If $\sum_{i=1}^{N} (D'_i - D_i) > 0$, the generated stream is infeasible with the given production capacity level and has to be discarded. In our numerical study, we have encountered no such infeasible streams. We set unit holding cost rate $h=1$, unit production cost $c=0$ and vary unit warming cost $\omega \in \{0.55, 0.65, 0.75, 0.85, 0.95\}$. The cold setup cost is selected as a function of the mean demand rate, $K = [J^2/2]\mu$ where J may be viewed as a *proxy* for the average length of production lot; we have $J \in \{2, 5\}$. The warm process threshold also varies as a function of the mean demand rate, $Q = \alpha\mu$ with $\alpha \in \{1, 1.3, 1.5, 2, 2.3, 3.5\}$. The tightness of capacity is attained by selecting the imposed physical capacity as $R = (\gamma K/\omega) + Q$ with $\gamma \in \{0.3, 0.4, 0.5, 0.6, 0.7, 0.8, 0.9, 1.0\}$. Note that γ corresponds to the ratio of keeping the process warm for one period and a cold setup.

We considered the problem of obtaining a production plan (i) statically (when demands for the entire problem horizon are known at the beginning of the problem horizon) and (ii) on a rolling horizon basis with given planning horizon lengths (when demands are revealed sequentially). We used 19 different planning horizon lengths *PHL* (in terms of number of periods);

Table 1
Summary of heuristics.

Heuristic	*Construction basis*	*Stopping rule with* $\hat{n}$
#1	Silver–Meal	$\frac{K+C_{u,u+n}}{n} < \frac{K+C_{u,u+n+1}}{n+1}$ $\hat{n}=1+\max\left\{n : \frac{K+C_{u,u+n}}{n} \geq \frac{K+C_{u,u+n+1}}{n+1}\right\}$
#2	Part period balancing	$C_{u,u+n} > K$ $\hat{n}=\max\{n : C_{u,u+n} \leq K\}$
#3	Least Unit Cost	$\frac{K+C_{u,u+n}}{\sum_{i=u}^{u+n-1} D_i} < \frac{K+C_{u,u+n+1}}{\sum_{i=u}^{u+n} D_i}$ $\hat{n}=1+\max\left\{n : \frac{K+C_{u,u+n}}{\sum_{i=u}^{u+n-1} D_i} \geq \frac{K+C_{u,u+n+1}}{\sum_{i=u}^{u+n} D_i}\right\}$
#4	Least Total Cost	$C_{u,u+n} > K$ $\hat{n}=\arg\min_{n \in \{\tilde{n},\tilde{n}+1\}} \lvert K-C_{u,u+n}\rvert,$ $\tilde{n}=\max\{i : C_{u,u+i} \leq K\}$
#5	Economic production quantity	$n > \widehat{q^*}/\overline{D}$ $\hat{n}=\lfloor \widehat{q^*}/\overline{D} \rfloor$
#6	Groff	$\frac{K+\tilde{C}_{u,u+n}}{n/N} < \frac{K+\tilde{C}_{u,u+n+1}}{(n+1)/N}$ $\hat{n}=1+\max\left\{n : n(n+1)D_{u+n} < \frac{[\widehat{q^*}]^2}{\overline{D}}\right\}$
#7	McLauren's Order Moment	$K+C_{u,u+n} > \widehat{CC^*}$ $\hat{n}=\max\{n : K+C_{u,u+n} \leq \widehat{CC^*}\}$
#8	McLauren's Order Moment 2	$\sum_{i=u}^{u+n-1} D_i > \widehat{q^*}$ $\hat{n}=\max\{n : \sum_{i=u}^{u+n-1} D_i \leq \widehat{q^*}\}$
#9	Wagner–Whitin	$\breve{x}_{u+n} > 0$ $\hat{n}=\min\{n : \breve{x}_{u+n} > 0,\ n > 1\}-1$

$PHL \in \{2,3,\ldots,19,20\}$. Note that the static solution may be viewed as having a planning horizon equal to the problem horizon.

Overall, our experimental set contains 480 ($=5\times2\times6\times8$) *parameter instances* for each of the four levels of demand variance and a given planning horizon length. (Capacity and warm threshold values have been truncated if non-integer for consistency with demands.) The average fraction of *zero* demand values in the generated replication streams is 0%, 2.32% and 15.72% for $\sigma/\mu=0.1$, 0.5 and 1 respectively. Across all planning horizon lengths, we have 9600 ($=480\times19+480$) *experiment instances* for each of the four levels of demand. For the static case, an experiment instance and a parameter instance coincide by definition. For each particular experiment instance – a particular combination of system parameter values – we have generated 30 demand stream *replications* for non-zero demand variances yielding a total of 873,600 ($=480\times(1+19)\times30\times3$ plus $480\times(1+19)$ for $\sigma=0$) *problem instances* with 43,680 instances for the static case and the rest for the rolling horizon case.

We discuss our findings for the static and rolling horizon cases separately; we begin with the former.

5.1. Static case

For the static case, we use as the benchmark the optimal solution to problem (P) with the entire problem horizon. The total cost over the horizon under a particular heuristic THC_i and the optimal total cost over the horizon THC_{opt} are computed as the average values across the replications for an experiment instance. For computing the total cost under a particular heuristic for a problem instance, we used the algorithm employing *Program_FindProductionPlan* (i',N') directly by setting $i'=1$ and $N'=N$. For each experiment instance, we measure heuristic performances in terms of percentage deviations from the optimal total cost which is computed for Heuristic $i \in \{\#1,\ldots,\#9\}$ as follows:

$$\Delta_i\% = \frac{THC_i - THC_{opt}}{THC_{opt}} \times 100$$

Thus, for each heuristic, we obtain a distribution of percentage deviations for four different demand variance levels based on replication-averages over all 480 experiment instances. Table 2 provides the average and the five number summary (maximum, third, second and first quantiles and the minimum) of the deviations. As σ increases, percentage deviations also increase for all heuristics. All heuristics have left-skewed performance distributions for all demand variance values. The performances of *Heuristics* #3, #4 and #5 are more sensitive – in that order – to σ. When we consider the performances in the average and median percentage deviations from the optimal. *Heuristic* #1 performs best in both measures. The ranking of other heuristics changes with the respect to the variance in demand and the performance measure. *Heuristic* #9 ranks very low when demand variance is low; but its ranking improves with the increase in σ. In terms of average percentage deviation, *Heuristics* #6, and #9 perform closely; similar behavior is observed with respect to their median performance as σ increases.

In Table 3, we tabulate the fraction of problem instances (expressed in % points) in which a particular heuristic strictly dominates another one in a pairwise fashion. For example, *Heuristic* #1 strictly dominates *Heuristic* #2 in 79.1% of all 43,680 problem instances, strictly dominates *Heuristic* #3 in 98.4% and so forth. An instance in which a heuristic dominates another does not necessarily imply that the dominating heuristic gives the best heuristic solution for that instance. There is no one heuristic that dominates nor is dominated by all others in all instances but there are clear winners. *Heuristic* #1 is the best and *Heuristic* #3 is the worst in this ranking although, in mean and median deviation percentages, it performs well for low demand variances.

Table 2
The average and five number summary (maximum, third, second, first quantiles, minimum) of deviations from the benchmark (averaged over 30 replications and 19 planning horizons). The benchmark for the static case is the optimal solution; the benchmark for the rolling horizon case is the rolling horizon DP solution.

		$\sigma/\mu=0$	$\sigma/\mu=0.2$
(for $\sigma/\mu=0$ and $\sigma/\mu=0.2$)			
Static	#1	12.6;(50.1;14.8;2.7;0.2;0.0)	16.1;(53.2;24.1;5.4;1.8;1.0)
	#2	15.9;(83.3;25.2;4.3;0.3;0.0)	19.7;(82.8;32.2;5.9;2.0;1.5)
	#3	12.7;(50.1;14.8;2.7;0.2;0.0)	18.1;(59.0;29.7;9.0;2.6;1.7)
	#4	22.0;(96.4;38.1;9.0;0.3;0.0)	26.6;(98.8;43.7;9.1;6.0;4.1)
	#5	22.8;(134.6;35.1;4.7;0.2;0.0)	27.4; 132.7;42.2;16.1;1.9;1.5)
	#6	22.6;(134.6;35.1;4.5;0.2;0.0)	26.7;(133.9;43.1;9.5;3.3;1.3)
	#7	28.2;(135.2;39.9;11.5;9.0;0.0)	31.8;(138.6;43.4;18.6;10.4;5.0)
	#8	22.7;(134.6;35.1;4.7;0.2;0.0)	26.7;(132.6;42.3;14.5;2.1;1.5)
	#9	23.5;(134.6;35.1;9.4;0.2;0.0)	26.5;(132.1;40.4;12.8;1.4;0.1)
Rolling horizon	#1	10.8;(50.0;11.9;0.0;0.0; −8.9)	14.5;(53.3;23.0;3.8;1.0; −2.7)
	#2	14.2;(83.3;24.6;0.2;0.0;0.0)	18.9;(82.7;31.0;6.9;1.4; −2.3)
	#3	10.8;(50.0;11.9;0.0;0.0; −8.9)	16.5;(59.2;27.9;7.4;1.7; −1.9)
	#4	22.7;(96.3;35.0;11.1;5.7; −8.9)	25.6;(98.5;41.6;10.2;5.5;0.2)
	#5	21.3;(134.7;33.0;4.7;0.0; −8.8)	26.2;131.6;39.2;14.8;1.4; −2.2)
	#6	21.1;(134.7;33.0;4.4;0.0; −8.9)	25.9;(133.6;40.0;12.6;2.8; −2.8)
	#7	21.4;(134.7;33.0;4.7;0.0; −8.8)	26.3;(128.0;38.5;14.4;3.7; −2.1)
	#8	21.3;(134.7;33.0;2.2;0.0; −8.8)	25.1;(132.3;39.7;11.9;1.6; −2.3)
	#9	20.7;(134.7;35.2;1.4;0.0;0.0)	25.2;(131.7;38.1;12.1;1.2;0.0)
(for $\sigma/\mu=0.5$ and $\sigma/\mu=1$)			
		$\sigma/\mu=0.5$	$\sigma/\mu=1$
Static	#1	19.0;(56.8;29.0;8.4;5.5;3.0)	21.4;(54.5;36.7;11.2;8.1;4.1)
	#2	23.7;(81.1;38.2;10.6;7.8;3.9)	26.1;(71.2;42.0;14.7;10.6;5.5)
	#3	29.6;(67.4;38.0;27.3;15.4;12.3)	40.7;72.1;47.9;39.2;33.6;21.5)
	#4	33.4;(104.2;49.2;18.7;14.8;11.1)	48.0;107.7;56.8;39.8;32.0;29.4)
	#5	36.3;(134.8;51.8;25.1;11.4;9.9)	43.0;121.9;51.6;33.0;24.6;21.6)
	#6	29.3;(131.0;43.8;11.4;6.4;2.8)	29.0;(111.5;43.5;12.4;7.7;4.2)
	#7	38.3;(139.1;55.1;22.0;15.6;11.8)	42.2;131.2;52.0;29.6;23.2;16.7)
	#8	31.2;(132.6;43.3;19.9;7.8;3.9)	30.7;(119.8;45.6;16.3;8.8;4.9)
	#9	28.7;(132.3;43.3;11.8;5.4;0.2)	28.5;(117.6;43.3;12.2;6.8;1.6)
Rolling horizon	#1	17.6;(56.7;28.9;6.9;4.0;0.5)	20.0;(54.4;36.2;9.9;6.5;2.1)
	#2	23.0;(81.0;35.8;11.4;7.6;1.7)	24.7;(71.3;41.4;13.1;9.0;4.3)
	#3	27.5;(67.5;35.0;24.0;12.9;7.3)	37.1;(62.1;45.0;35.7;29.6;16.5)
	#4	30.7;(102.1;45.9;16.7;12.4;6.6)	34.8;94.3;47.3;23.5;19.0;13.7)
	#5	37.2;(130.3;55.4;24.6;11.4;7.0)	60.4;398.6;66.3;39.2;25.8;17.0)
	#6	30.1;(132.4;45.7;16.8;6.6;0.7)	34.7;(189.3;52.2;24.4;9.6;3.8)
	#7	35.8;(134.4;53.2;22.6;11.3;5.1)	50.4;393.9;58.5;33.0;16.8;10.6)
	#8	30.2;(132.3;43.7;17.9;7.2;1.7)	33.9;(119.7;51.9;25.1;9.5;4.1)
	#9	27.8;(131.8;42.9;12.4;5.3;0.0)	27.9;(117.6;42.0;12.6;6.6;1.2)

Next, we present results for each heuristic on the average performance of a particular heuristic *vis-a-vis* the optimal solution in the instances where it is the best heuristic. The entry for row *j* in Table 4 shows, respectively, the average percentage deviation of heuristic *j* and the average percentage deviation of all heuristics from the optimal in the instances when the heuristic is the best (or among the best) and the percentage of such instances in parentheses. For example, for $\sigma/\mu=0$, *Heuristic* #1 is the best performer (or among the best performers) in 81.7% of the experiment instances with an average deviation from the optimal of 11.3% while the average deviation for all heuristics is 20.8%. If a heuristic has no entries (*e.g.*, *Heuristic* #4), it implies that the heuristic has never performed best for that category of demand. *Heuristics* #1 and #9 turn out to be best performers for a large fraction of experiment instances in the overall, followed by *Heuristic* #3. In Table 4, we also tabulate the percentage of problem instances in which a particular heuristic has been found to be the best heuristic, and the percentage of instances in which it is within 2% and 5% proximity of the best heuristic performance. For example, *Heuristic* #1 results in the best cost performance in 42.9% of the problem instances, is within 2% of the best cost performance in 80.8% of the instances, within 5% of the best cost performance in 98.2% of the instances. *Heuristic* #1 ranks first and *Heuristic* #7 ranks last by all performance criteria.

Table 3
Fraction of problem instances (in % points) of dominance in pairwise comparison.

	Heuristic	#1	#2	#3	#4	#5	#6	#7	#8	#9
Static	#1	–	79.1	**98.4**	84.6	59.9	92.3	81.6	96.0	81.4
	#2	20.5	–	**93.1**	59.8	33.1	76.6	59.0	81.4	47.8
	#3	1.4	6.7	–	11.3	8.3	41.6	41.0	**49.4**	9.1
	#4	14.9	39.2	**88.6**	–	24.7	79.3	58.4	78.0	37.5
	#5	39.6	65.9	**91.7**	74.3	–	83.8	63.7	80.1	59.6
	#6	7.1	22.4	58.4	17.8	15.2	–	52.2	**68.1**	11.7
	#7	12.9	40.5	58.9	40.9	35.9	45.8	–	**61.2**	32.1
	#8	3.6	18.1	**50.4**	21.6	19.5	31.2	38.3	–	8.2
	#9	15.4	51.8	**90.8**	54.8	39.9	84.5	62.6	90.3	–
Rolling horizon	#1	–	87.5	80.5	**97.7**	94.5	92.2	96.9	93.1	72.0
	#2	9.3	–	60.7	87.0	84.7	59.6	**87.6**	56.9	44.5
	#3	13.3	33.9	–	54.1	53.6	45.9	**55.9**	43.6	40.2
	#4	2.1	11.7	45.7	–	**63.2**	28.9	46.4	29.5	25.8
	#5	4.9	11.2	**44.2**	36.2	–	25.2	31.3	21.2	20.3
	#6	7.2	39.9	53.6	70.7	**73.8**	–	67.7	55.2	32.9
	#7	2.6	11.8	43.6	53.3	**67.3**	31.3	–	28.9	27.8
	#8	6.3	31.0	55.6	70.2	**76.6**	43.8	70.0	–	28.2
	#9	27.5	54.9	59.3	74.0	**79.0**	66.5	71.6	71.2	–

Finally, we attempt to identify the range of experiment parameter values where each particular heuristic is the best or among the best. Identification of such operating environment

Table 4
Mean performances of individual heuristics and all heuristics in the experiment instances in which a particular heuristic is best or among the best and fraction of problem instances (in % points) in which a particular heuristic is the best (or among the best) or Whitin 2% and 5% of the best heuristic solution for the static case. (* corresponds to 3 instances out of 14,400.)

	$\sigma/\mu=0$	$\sigma/\mu=0.2$	$\sigma/\mu=0.5$	$\sigma/\mu=1$	*Overall*	$\leq 2\%$	$\leq 5\%$
#1	11.3/20.8 (81.7)	24.0/38.3 (47.8)	26.3/40.5 (48.4)	31.2/47.0 (31.2)	26.3/40.9 (42.9)	(80.8)	(98.2)
#2	8.4/16.8 (66.7)	22.9/42.6 (14.5)	16.6/29.8 (3.4)	18.2/30.3 (0.7)	20.2/37.3 (6.9)	(31.5)	(70.0)
#3	10.9/20.1 (82.3)	28.5/48.9 (20.6)	30.6/49.9 (10.0)	35.8/57.3 (15.3)	30.2/50.1 (16.0)	(35.4)	(49.0)
#4	3.8/7.2 (45.6)	15.4/21.5 (4.6)	–/–	–/–	12.5/17.9 (2.0)	(5.3)	(15.7)
#5	0.5/2.4 (46.7)	6.6/9.4 (2.8)	–/–	–/–	4.4/6.9 (1.4)	(14.1)	(23.4)
#6	0.4/2.4 (47.7)	5.8/9.7 (4.4)	7.0/16.3(5.6)	9.4/24.9 (6.7)	7.0/16.6 (6.0)	(34.3)	(67.3)
#7	6.0/10.5 (14.2)	39.6/43.4 (1.0)	34.5/36.6 (0.0*)	–/–	28.7/32.7 (0.5)	(2.2)	(12.3)
#8	0.5/2.4 (46.9)	16.0/18.7 (2.1)	21.7/26.2 (0.3)	8.8/22.1 (1.5)	9.9/ 15.4 (1.8)	(25.5)	(61.2)
#9	0.6/2.1 (43.8)	0.9/3.8 (27.9)	4.7/10.6 (32.6)	8.4/19.6 (44.6)	5.2/12.5 (35.1)	(53.8)	(65.7)

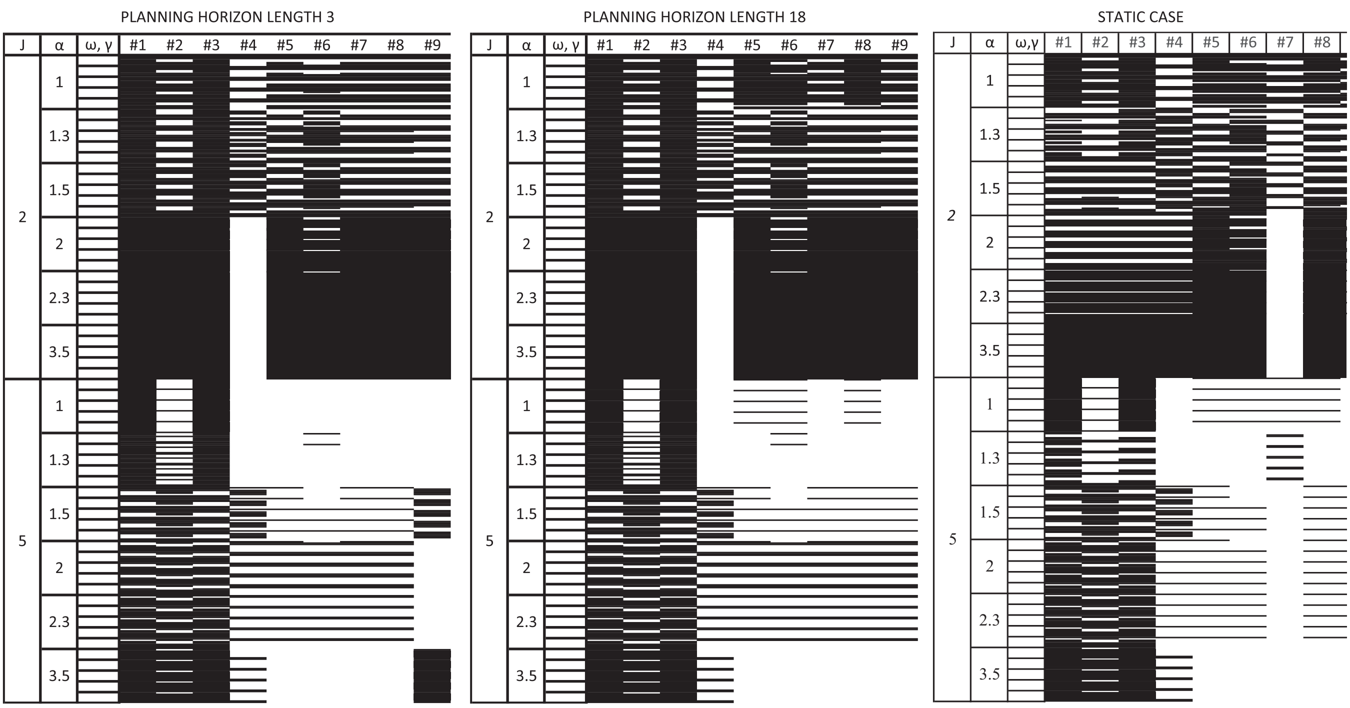

Fig. 1. Illustration of the experiment instances for which a heuristic is the best performer (or among the best), $\sigma/\mu=0$.

characteristics would help practitioners choose *a priori* the best heuristic for a given operational setting. In Figs. 1–4, we illustrate (as blackened) the regions of the experiment set for which a particular heuristic is the best (or among the best). Based on this illustration, we can make the following observations. For $\sigma=0$, for $J=2$, we see that all heuristics except #7 show comparable performance. When $J=5$, however, a delineation among the heuristics starts to emerge. *Heuristics* #1, #2 and #3 dominate the rest but, we see that there are regions for all heuristics in which they are the best (or among the best). (See Fig. 1.) As demand variability increases ($\sigma/\mu=0.2$), the deviation increases further. For low *J*, *Heuristics* #1 and #9 almost complement each other; that is, one is the best in the regions where the other is not. For high values of α, *Heuristic* #9 dominates. We see fewer instances where *Heuristics* #2, #4, #5, and #6 are the best (or among the best) performers. *Heuristics* #3, #7 and #8 are never among the best performers in this region. For high *J*, *Heuristic* #1 is the best for the large majority of experiment instances (especially for low α) followed by *Heuristics* #2, #3 and #4. For low and high ends of α values, *Heuristic* #1 dominates the rest of the heuristics. There are also instances where *Heuristics* #6, #7 and #8 are among the best performers; but *Heuristics* #5 and #9 are never the best in this region. (See Fig. 2.) For $\sigma/\mu=0.5$ and low *J*, *Heuristics* #1 and #9 are the best for most of the instances followed by only *Heuristic* #6; the rest is never among the best performers. We observe the complementary behavior of *Heuristics* #1 and #9 in this case as well. For low α, *Heuristic* #1 is the best whereas *Heuristic* #9 dominates for high α values. For high *J*, *Heuristic* #1 dominates clearly except for moderate values of α and high values of γ. It is complemented in those regions by *Heuristic* #3. *Heuristic* #9 starts to show presence for large α and γ. (See Fig. 3.) For high demand variance ($\sigma/\mu=1$), *Heuristic* #9 dominates for low *J* and high α values at the expense of *Heuristic* #1. *Heuristic* #6 also emerges as the best heuristic complementing *Heuristic* #9. For high *J* and low α, *Heuristic* #1 is the sole best heuristic but *Heuristics* #3 and #9 dominate it in the rest of the experiment instances. For moderate α, *Heuristic* #3 is more likely to be the better than *Heuristic* #9. (See Fig. 4.)

Next, we discuss our findings for the rolling horizon case.

5.2. Rolling horizon case

The rolling horizon case corresponds to the problem of obtaining the production plan dynamically as time progresses

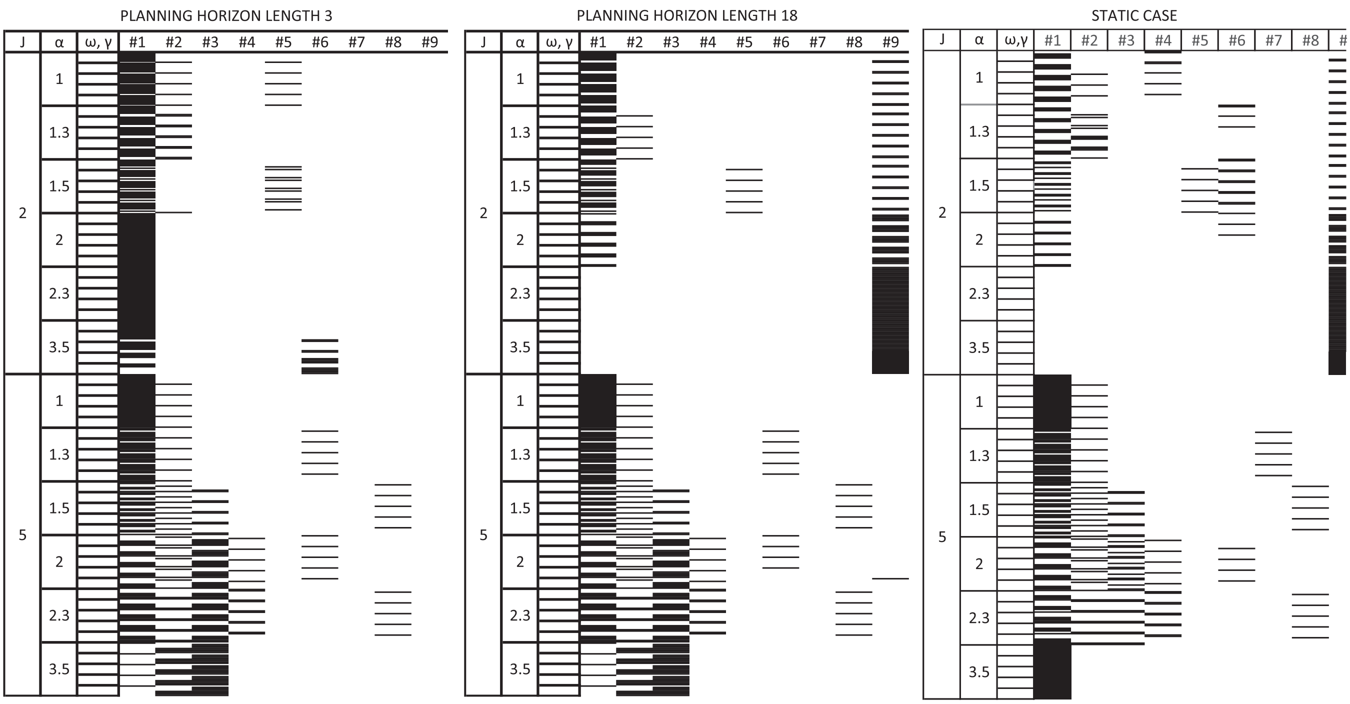

Fig. 2. Illustration of the experiment instances for which a heuristic is the best performer (or among the best), $\sigma/\mu = 0.2$.

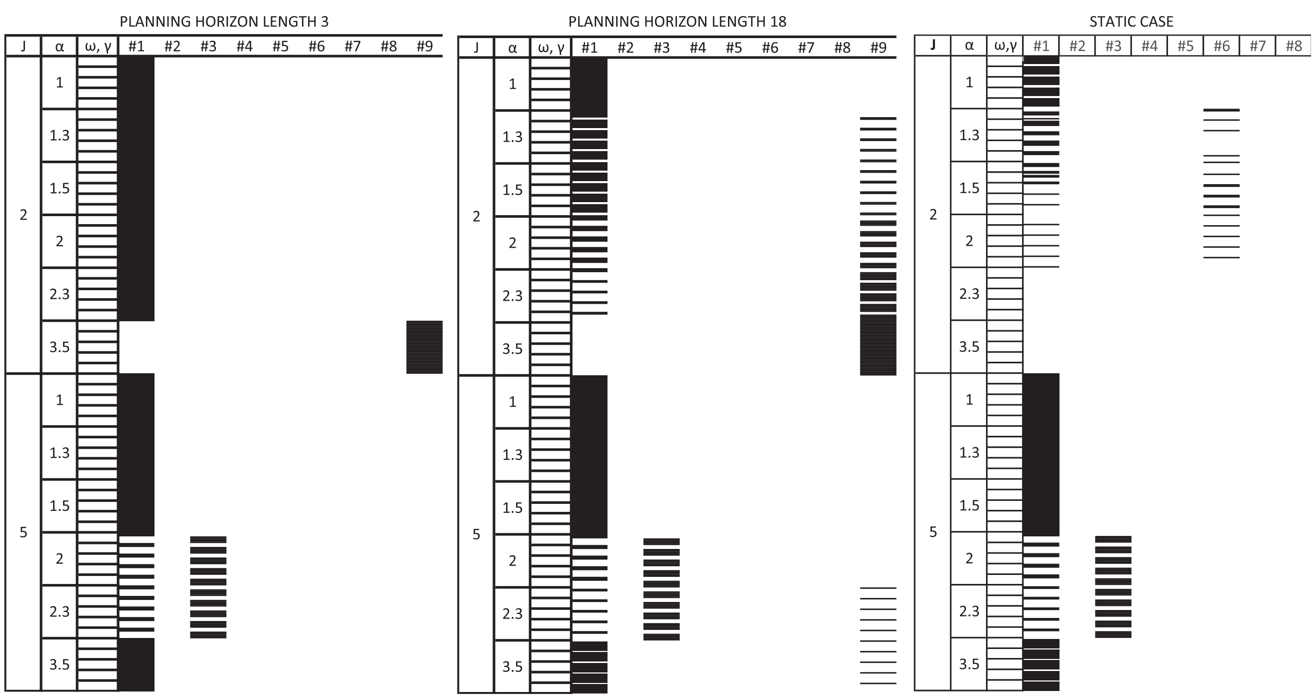

Fig. 3. Illustration of the experiment instances for which a heuristic is the best performer (or among the best), $\sigma/\mu = 0.5$.

and demands within a planning horizon are revealed sequentially. (That is, demands become known within a planning horizon as the horizon rolls; therefore, the deterministic demand condition for Proposition 1 hold for the new horizon length.) For this case, we use as the benchmark the dynamic programming (DP) solution of the problem obtained on a rolling basis with the given planning horizon; that is, we solved problem (*P*) sequentially as the problem horizon extended over time until we reached period N as the end of the planning horizon. The total cost over the entire problem obtained thus is denoted by THC_{DPS} and is computed as the average value across 30 replications (and nineteen different planning horizons for a parameter instance when so noted). We note that the benchmark rolling horizon DP solution may not give the best cost; this is illustrated by our results, especially for short planning horizon lengths. (We discuss this finding later below.) The total cost over the horizon under a particular heuristic THC_i is

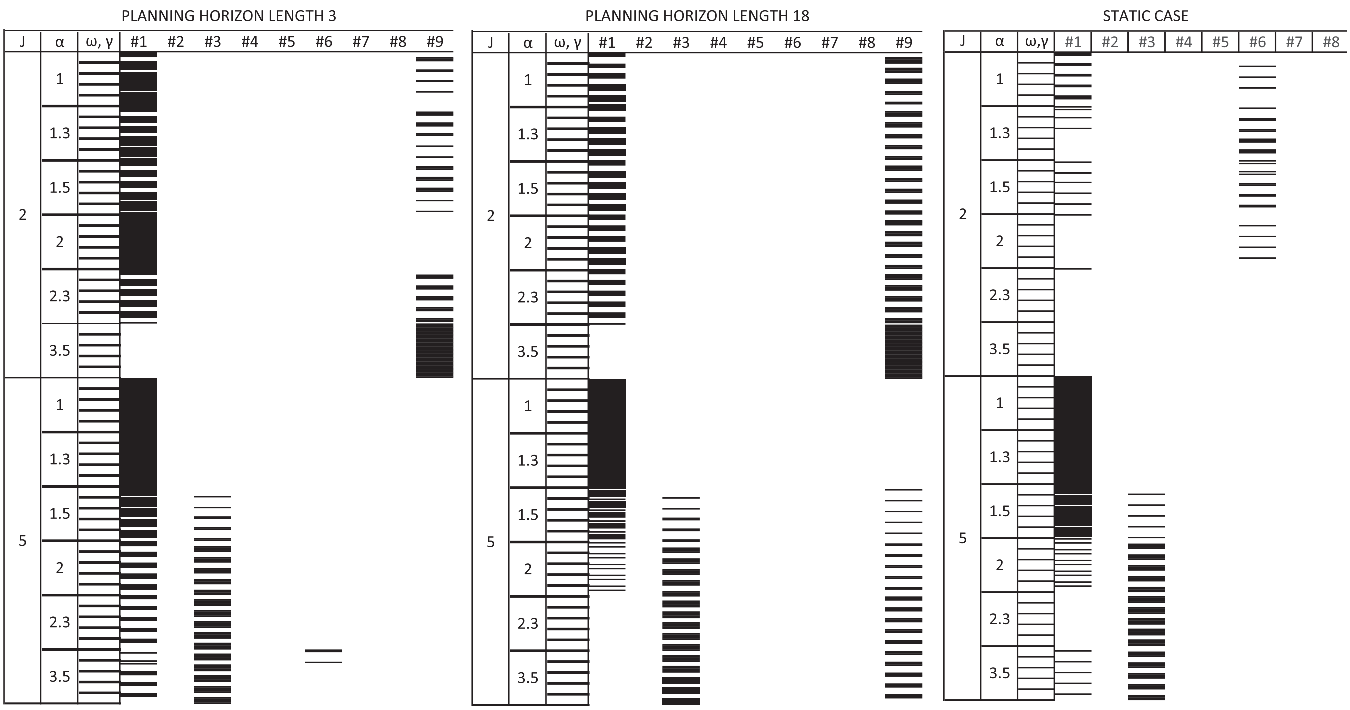

Fig. 4. Illustration of the experiment instances for which a heuristic is the best performer (or among the best), $\sigma/\mu=1$.

Table 5
Mean performances of individual heuristics and all heuristics in the experiment instances in which a particular heuristic is best or among the best and fraction of problem instances (in % points) in which a particular heuristic is the best (or among the best) or Whitin 2% and 5% of the best heuristic solution for the rolling horizon case.

	$\sigma/\mu=0$	$\sigma/\mu=0.2$	$\sigma/\mu=0.5$	$\sigma/\mu=1$	*Overall*	$\le 2\%$	$\le 5\%$
#1	9.5/18.6 (93.8)	19.1/32.4 (55.3)	20.1/34.8 (67.1)	23.1/43.6 (51.5)	16.8/58.0 (66.9)	(92.1)	(99.4)
#2	7.4/16.2 (72.1)	21.6/46.4 (14.1)	–/–	–/–	9.7/0.0 (21.5)	(31.4)	(68.5)
#3	9.5/19.9 (93.8)	19.2/39.7 (19.8)	29.9/51.4 (10.4)	31.4/53.2 (16.3)	14.9/87.7 (35.1)	(37.1)	(50.7)
#4	5.9/10.6 (27.1)	13.0/15.6 (3.1)	–/–	–/–	6.6/0.0 (7.6)	(3.9)	(17.6)
#5	0.8/2.3 (49.9)	0.8/2.4 (1.6)	–/–	–/–	0.8/0.0 (12.9)	(14.4)	(23.4)
#6	0.5/2.3 (51.0)	20.0/22.1 (2.4)	–/–	9.1/15.3(0.2)	1.4/1.0 (13.4)	(26.7)	(58.7)
#7	0.8/2.0 (47.9)	11.7/13.6 (0.1)	–/–	–/–	0.8/ 0.0 (12.0)	(8.8)	(27.0)
#8	0.8/2.3 (49.9)	25.8/28.3 (2.1)	–/–	–/–	1.8/ 0.0 (13.0)	(22.0)	(53.1)
#9	1.1/3.2 (46.0)	0.6/2.7 (21.5)	2.8/8.9 (22.5)	7.5/25.0 32.1)	3.0/ 4.2 (30.5)	(47.1)	(62.1)

computed also similarly as an average. To compute the total cost under a particular heuristic for a problem instance, we used the algorithm whose pseudo-code is given below.

begin
 $j := 1$
 $\hat{k} := 1$
 while $(j \le N - PHL + 1)$
  $i' := \hat{k}$
  $N' := j + PHL - 1$
  call Program_FindProductionPlan (i', N')
  $j := j + 1$
  $\hat{k} :=$ *last cold setup such that* $\hat{k} \le j$
end

The heuristic performance for each parameter instance for Heuristic $i \in \{\#1, \ldots, \#9\}$ is found as follows:

$$\Delta_i\% = \frac{THC_i - THC_{DPS}}{THC_{DPS}} \times 100$$

The analysis and discussion of our findings follow closely that for the static case above.

Table 2 provides the average and the five number summary (maximum, third, second and first quantiles and the minimum) of the deviations. As σ increases, percentage deviations also increase for all heuristics. All heuristics have left-skewed performance distributions for all demand variance values. The performances of *Heuristics* #3, #5 and #7 are more sensitive – in that order – to σ with maximum deviations under *Heuristics* #5 and #7 increasing while that of *Heuristic* #3 decreases. Across all variance levels, *Heuristic* #1 performs best in terms of both the average and the median. The ranking of other heuristics changes with the respect to the variance in demand and the performance measure. For deterministic demand, there are multiple heuristics resulting in the same performance. As demand variance increases, the difference among average percentage deviations also increases. Note that for deterministic demand *Heuristics* #1 and #2 have zero values for the median percentage deviation.

In Table 3, we tabulate the fraction of problem instances (expressed in % points) in which a particular heuristic dominates another one in a pairwise fashion. *Heuristic* #1 is the best and *Heuristic* #5 is the worst in this ranking.

Next, we present results for each heuristic on the average performance of a particular heuristic *vis a vis* the optimal solution in the instances where it is the best heuristic. The entry for column *j* in Table 5 shows, respectively, the average percentage deviation of heuristic *j* and the average percentage deviation of all heuristics from the optimal in the instances when the heuristic is the best (or among the best) and the percentage of such instances in parentheses. *Heuristic* #1 is the best performer overall when we compare the fraction of dominated instances and the deviation percentages. In Table 5, we also tabulate the percentage of problem instances in which a particular heuristic has been found to be the best heuristic, and the percentage of instances in which it is within 2% and 5% proximity of the best heuristic performance. *Heuristic* #1 ranks first and *Heuristics* #4, #5 and #7 are worst performers by all criteria.

Next, as in the static case, we provide, in Figs. 1–4, an illustration (as blackened) of the regions of the parameter set for which a particular heuristic is the best (or among the best) for one short and one longer planning horizon (*PHL*=3 and 18). We discuss the short planning horizon case first. For $\sigma=0$, for $J=2$, we see that all heuristics show comparable performance except *Heuristics* #4 for moderate to large α. When $J=5$, however, a delineation among the heuristics starts to emerge. *Heuristics* #1, #2 and #3 dominate the rest but, we see that there are regions for all heuristics in which they are the best (or among the best). (See Fig. 1.) As demand variability increases ($\sigma/\mu=0.2$) in Fig. 2, certain heuristics start to emerge as the sole best. Specifically, *Heuristic* #1 is the best performer over a vast majority of parameter instances. For large *J* and α, *Heuristics* #2 and especially #3 become better. *Heuristics* #7 and #9 are never the best; the rest perform well sporadically. For $\sigma/\mu=0.5$, (in Fig. 3) *Heuristic* #1 is clearly the best performer with *Heuristic* #3 for some instances with large *J* and moderate α, and *Heuristic* #9 only for low *J* and large α. The rest of the heuristics is never the best. As demand variance increases further ($\sigma/\mu=1$) in Fig. 4, *Heuristic* #1 deteriorates and *Heuristics* #9 and #3 replace it. *Heuristic* #6 performs well for few instances of large *J* and moderate α. For longer planning horizons, there is not a discernable difference compared to short horizon for $\sigma=0$; for large *J*, *Heuristic* #9 deteriorates for large α but improves slightly for moderate α. For moderate α and large *J*, *Heuristics* #5, #6, and #8 start to perform well. As demand variance increases, *Heuristics* #1, #3 and #9 start to dominate the rest with *Heuristic* #3 being best for large *J* and moderate α, and #9 for small *J* and moderate to high α.

Next, we present our observations on the impact of planning horizon length on heuristic performance. In terms of best performing heuristics, there is a clear separation among *Heuristics* #1 and #3 and the rest. The performances deteriorate but stabilize as planning horizon extends. This is to be expected: The benchmark solution approaches the optimal static solution for longer horizon lengths but, as the planning horizon gets very short, it becomes more suboptimal as evidenced by better solutions obtained via *Heuristic* #1 for such instances.

For $\sigma/\mu=0$ and low *J*, *Heuristics* #1 and #3 are the best performers over all planning horizon lengths, followed by *Heuristics* #6, #5 and #8. *Heuristic* #4 performs distinctly badly whereas the remaining heuristics form a cluster with higher deviations. For large *J*, the performance deteriorates for all heuristics. *Heuristics* #1 and #3 are again the best performers followed by *Heuristics* #2, #6 and then by *Heuristics* #4 and #9. The rest form a separate grouping with very close performances. Overall, the heuristic performance is similar to that for large *J* since the deviations are larger for large *J*. For $\sigma/\mu=0.2$ and low *J*, *Heuristic* #1 is clearly the best performer over all planning horizon lengths, followed by a grouping comprising *Heuristics* #8, #2 and #3. Note that for certain short planning horizons, the deviations of *Heuristic* #1 from the rolling horizon DP solution are negative which indicates that, on average, *Heuristic* #1 provides better solutions. (We elaborate on this aspect later.) The second grouping is followed by *Heuristics* #9 and #6. Heuristic #9 improves over longer planning horizon lengths whereas *Heuristic* #6 deteriorates. The rest follow the trend of the second grouping but with larger deviations. For large *J*, the performance deteriorates for all heuristics. *Heuristic* #1 is again the best performer followed closely by *Heuristic* #3, and then by *Heuristics* #2 and #6. The rest is in a distinctly separate grouping. Overall, *Heuristic* #1 performs the best followed by *Heuristics* #3 and #2. The remaining heuristics form a distinct grouping with similar behavior. For $\sigma/\mu=0.5$ and low *J*, *Heuristic* #1 is the best performer over all planning horizon lengths, followed by *Heuristics* #9 and #2.

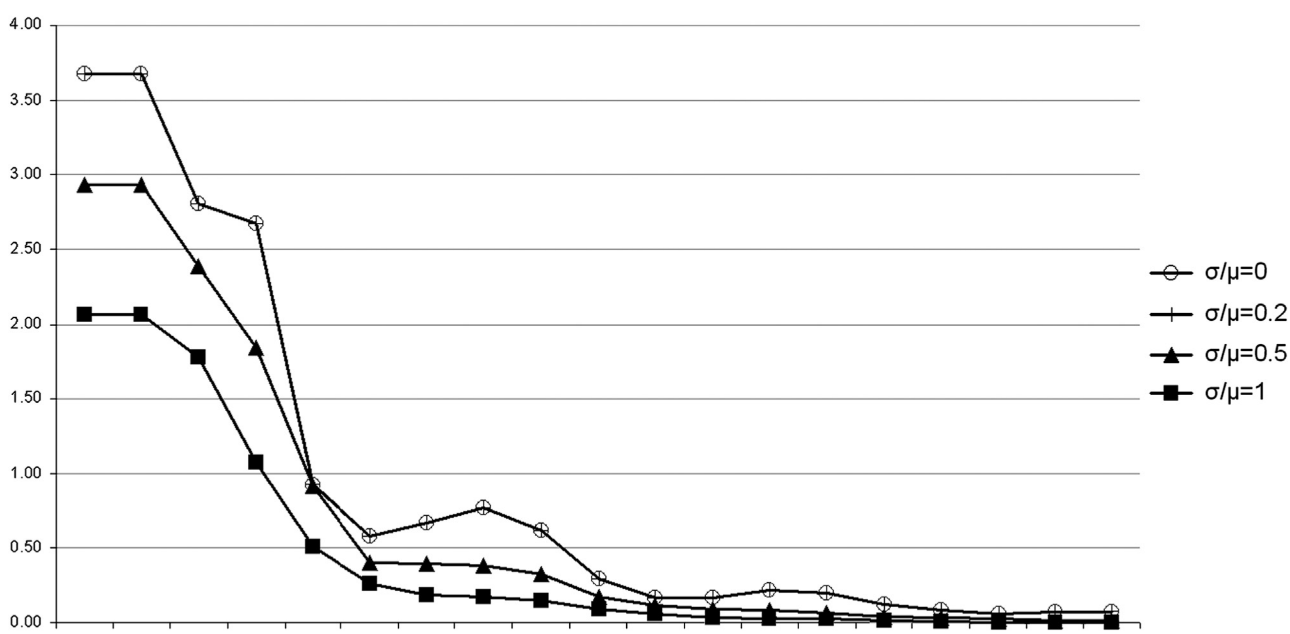

Fig. 5. The average percentage deviation of the rolling DP solution from the static case versus planning horizon lengths.

Heuristics #6 and #8 deteriorate over longer planning horizon lengths; the rest is relatively stable. For large *J*, the performance deteriorates for all heuristics. *Heuristic*#1 is again the best performer followed by *Heuristics* #3 and #2. The rest is very close to each other. Overall, *Heuristic* #1 performs the best followed by *Heuristics* #2, #3 and #9. *Heuristics* #6, #8, #4 and #5, #7 form the remaining groupings. For $\sigma/\mu = 1$ and low *J*, *Heuristics* #1 and #9 have very similar performances and are the best performers over all planning horizon lengths, followed closely by *Heuristic #2*. *Heuristics* #6 and #8 deteriorate over longer planning horizon lengths. *Heuristics* #3 and #4 are stable throughout but *Heuristic* #7 deteriorates fast stabilizing over longer horizon lengths. *Heuristic* #5 has the worst performance throughout. For large *J*, the performance deteriorates for all heuristics. *Heuristic* #1 is again the best performer followed by *Heuristics* #3 and #2. *Heuristics* #9, #6, #8, #7 and #4 exhibit similar performances with #5 being the worst. Overall, *Heuristic* #1 performs the best followed by *Heuristics* #2 and #9. *Heuristics* #6, #3, #4 and #8 behave similarly with *Heuristics* #7 and #5.

Finally, we consider the impact of planning horizon lengths on the DP solution *vis a vis* the static solution.

We first consider the impact of introducing a planning horizon and obtaining the production plan on a rolling horizon basis. In Fig. 5 we present the average percentage deviation of the rolling horizon DP solution from the static case. We observe that as the planning horizon extends, the rolling solution approaches that of the static case. Interestingly we observe the highest percentage deviation for $\alpha/\mu = 0$. For variable demand patterns, the percentage deviations increase as demand variance increases.

When we consider the percentage of problem instances in which a particular heuristic results in a lower total cost than the rolling horizon DP solution for a planning horizon length, we observed the largest fraction to be for $\alpha/\mu = 0.2$ reaching 45.2% for short horizon lengths. We see that as demand variance increases such instances become rare; we observed no such instances for $\alpha/\mu = 1$. However, there is not a monotone behavior. For the deterministic case, we observed dominated instances for all planning horizon lengths but reported a portion for brevity. For $\alpha/\mu = 0.2$, we observed no such instances for any heuristic beyond *PH*=11; for $\alpha/\mu = 0.5$, beyond *PH*=3. For $\alpha/\mu = 0$, the fraction is lower but still sizable for *Heuristics* #1, #3, #4 and #6. These observations lead us to the conclusion that use of heuristics may be beneficial in practice. Overall, the fraction of instances that the heuristics #1 through #9 dominates the rolling DP solution has been found as : 2.6, 1.7, 1.4, 0.1, 1.4, 1.1, 1.0, 1.6 and 0.1. In Table 6, we present the average percentage savings that one gets by using *Heuristics* #1 instead of the DP solution on a rolling horizon basis; we focus on *Heuristics* #1 because it turns out to be the best performer based on our entire numerical study. As expected, the length of a planning horizon *vis a vis* the proxy for average production run length, *J* is an important factor.

Table 6
The percentage deviation of *Heuristic* #1 from the DP solution on a rolling horizon basis across all problem instances for different planning horizon lengths and *J*.

J	*Δ%*	*PH*								
		2	3	4	5	6	7	8	9	10
2	Max	4.1	4.1	1.6	1.6	0.6	0.5	0.2	0.0	0.1
	Median	1.7	1.7	0.5	0.5	0.1	0.2	0.2	0.2	0.0
	Mean	1.8	1.8	0.5	0.6	0.2	0.2	0.2	0.2	0.0
5	Max	1.2	1.2	1.2	1.2	0.0	0.0	0.0	0.0	0.0
	Median	0.5	0.5	0.5	0.3	0.0	0.0	0.0	0.0	0.0
	Mean	0.5	0.5	0.5	0.4	0.0	0.0	0.0	0.0	0.0

6. Conclusion

In this work, we have proposed nine rule-based lot sizing heuristics for a warm/cold process (defined as the one which can be kept warm for the next period at an additional linear cost if the production quantity in the current period is at least a positive threshold amount). Due to the nature of the stopping rules, the proposed heuristics fall into two categories: quantity based and cost based. For quantity based heuristics, we use an adaptation of the EPQ model. The stopping rule determines the size of the production lots. For all heuristics, the production schedule (over possibly consecutive production periods) within a production lot is determined by the optimal results obtained for the warm/cold process which minimize the total costs. In a numerical study, we have examined the performance of the proposed heuristics. We find that, overall, EPQ-based heuristics are dominated by those constructed on the basis of costs. Our findings further indicate that there is not a single heuristic that is best for all parameter settings. In terms of total cost, *Heuristics* #1 and #9 perform best for the static case but for rolling horizon settings, *Heuristic* #1 is clearly the best. In terms of fraction of experiment instances where a particular heuristic dominates others, *Heuristic* #1 is the best followed by #9 and #3. In the numerical study, we have also identified operating environments for which the proposed heuristics would perform best. In general, but especially for large demand variability (σ/μ), the heuristics constructed via EPQ-based rules (*Heuristics* #5 through #8) perform badly. As the warm/cold process approaches the classical problem setting, *Heuristic* #9 starts to dominate *Heuristic* #1. This happens for small *J*, large α and large γ/ω (with resulting large *R*). We observe that *Heuristic* #9 is replaced by *Heuristic* #3 for relatively medium to large values of α. The intuition behind these performance behavior may be as follows. The heuristics are constructed in two steps: (i) determination of the production lot size (according to the specified stopping rule), and (ii) determination of the production schedule (according to Proposition 1). The second decision is always optimal by construct. Thus, the performance of the heuristics primarily depends on how well the lot size is determined. The Wagner–Whitin solution (resulting in *Heuristic* #9) is, by definition, optimal for the static classical problem; our findings indicate that this uncapacitated solution provides also very good approximations for the lot size in a warm/cold process. In the case of *Heuristic* #1, it is the only one that uses a stopping rule that is based on cost rate optimization. Its good performance indicates that this criterion results in good lot sizes for a warm/cold process, as well. A similar explanation may be valid for the performance of *Heuristic* #3 which is based on a stopping rule minimizing costs per unit.

An important result of our numerical studies is that, when used on a rolling horizon basis, a heuristic solution for a warm/cold process may also perform better costwise than a solution obtained using a dynamic programming approach especially for short planning horizons and small *J*. This finding is consistent with similar studies on the classical problem (Stadtler, 2000; Heuvel and Wagelmans, 2005). Hence, investigation and implementation of heuristics for warm/cold process settings may be economically beneficial in practice as well as important from a purely theoretical perspective.

References

Berk, E., Toy, A.Ö., Hazir, Ö., 2008. Single item lot-sizing problem for a warm/cold process with immediate lost sales. European Journal of Operational Research 187 (3), 1251–1267.
Brahimi, N., Dauzere-Peresm, S., Najid, N.M., Nordli, A., 2006. Single item lot sizing problems. European Journal of Operational Research 168, 1–16.

Buschkühl, L., Sahling, F., Helber, S., Tempelmeier, H., 2010. Dynamic capacitated lot-sizing problem: a classification and review of solution approaches. OR Spectrum 32, 231–261.

Chubanov, S., Kovalyov, M.Y., Pesch, E., 2008. A single-item economic lot-sizing problem with a non-uniform resource: approximation. European Journal of Operational Research 189, 877–889.

DeMatteis, J.J., 1968. An economic lot-sizing technique I: the part period algorithm. IBM Systems Journal 7 (1), 30–38.

Federgruen, A., Tzur, M., 1991. A simple forward algorithm to solve general dynamic lot sizing models with n periods in $O(n \log n)$ or $O(n)$ time. Management Science 37 (3), 909–925.

Feng, Y., Chen, S., Kumar, A., Lin, B., 2011. Solving single-product economic lot-sizing problem with non-increasing setup cost, constant capacity and convex inventory cost in $O(n \log n)$ time. Computers and Operations Research 38, 717–722.

Florian, M., Lenstra, J.K., Rinnooy Kan, H.G., 1980. Deterministic production planning: algorithms and complexity. Management Science 26 (7), 669–679.

Glock, C.H., 2010. The joint economic lot size problem: a review. International Journal Production Economics 135, 671–686.

Groff, G.K., 1979. A lot-sizing rule for time phased component demand. Production and Inventory Management 20 (1), 47–53.

Guner Goren, H., Tunali, S., Jans, R., 2010. A review of applications of genetic algorithms in lot sizing. Journal of Intelligent Manufacturing 21, 575–590.

Hardin, J.R., Nemhauser, G.L., Savelsbergh, M.W.P., 2007. Analysis of bounds for a capacitated lot-sizing problem. Computers and Operations Research 34, 1721–1743.

Heuvel, W., Wagelmans, A.P.M., 2005. A comparison of methods for lot-sizing in a rolling horizon environment. Operations Research Letters 33 (5), 486–496.

Heuvel, W., Wagelmans, A.P.M., 2006. An efficient dynamic programming algorithm for a special case of the capacitated lot-sizing problem. Computers and Operations Research 33, 3583–3599.

Jans, R., Degraeve, Z., 2007. Meta-heuristics for dynamic lot sizing: a review and comparison of solution approaches. European Journal of Operational Research 177 (3), 1855–1875.

Khouja, M., Mehrez, A., 1994. Economic production lot size model with variable production rate and imperfect quality. The Journal of the Operational Research Society 45 (12), 1405–1417.

Larsen, C., 1997. Using a variable production rate as a response mechanism in the economic production lot size model. The Journal of the Operational Research Society 48 (1), 97–99.

Manne, A.S., 1958. Programming of economic lot sizes. Management Science 4 (2), 115–135.

Narasimhan, S., McLeavy, D.W., 1995. Production Planning and Inventory Control, 2nd ed.. Prentice-Hall, Englewood Cliffs.

Narayanan, A., Robinson, P., 2010. Evaluation of joint replenishment lot-sizing procedures in rolling horizon planning systems. International Journal of Production Economics 127, 85–94.

Ng, C.T., Kovalyov, M.Y., Cheng, T.C.E., 2010. A simple FPTAS for a single-item capacitated economic lot-sizing problem with a monotone cost structure. European Journal of Operational Research 200, 621–624.

Pochet, Y., Wolsey, L.A., 2010. Single item lot-sizing with non-decreasing capacities. Mathematical Programming 121, 123–143.

Quadt, D., Kuhn, H., 2008. Capacitated lot-sizing with extensions: a review. Quarterly Journal of Operations Research 6, 61–83.

Robinson, E.P., Sahin, F., 2001. Economic production lot sizing with periodic costs and overtime. Decision Sciences 32 (3), 423–452.

Sahin, F., Robinson, P., Gao, L., 2008. Master production scheduling policy and rolling schedules in a two-stage make-to-order supply chain. International Journal of Production Economics 115 (2), 528–541.

Silver, E.A., Meal, H.A., 1973. A heuristic for selecting lot size requirements for the case of a deterministic time varying demand rate and discrete opportunities for replenishment. Production and Inventory Management 14 (2), 64–74.

Simpson, N.C., 2001. Questioning the relative virtues of dynamic lot sizing rules. Computers and Operations Research 28 (9), 899–914.

Stadtler, H., 2000. Improved rolling schedules for the dynamic single-level lot-sizing problem. Management Science 46, 318–326.

Staggemeier, A.T., Clark, A.R., 2001. A survey of lot-sizing and scheduling models. In: 23rd Annual Symposium of the Brazilian Operational Research Society (SOBRAPO), Brazil, pp. 938–947.

Toy, A.Ö., Berk, E., 2006. Dynamic lot sizing problem for a warm/cold process. IIE Transactions 38 (11), 1027–1044.

Vollmann, T.E., Berry, W.L., Whybark, D.C., 1997. Manufacturing Planning and Control Systems, 4th ed.. McGraw-Hill.

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.M., Whitin, T.M., 1958. Dynamic version of the economic lot size model. Management Science 5 (1), 89–96.