

A 2-Approximation Algorithm for Stochastic Inventory Control Models with Lost Sales

Retsef Levi

Sloan School of Management, MIT, Cambridge, MA, 02139, USA
email: retsef@mit.edu

Ganesh Janakiraman

IOMS-OM Group, Stern School of Business, New York University, New York, NY 10012, USA
email: gjanakir@stern.nyu.edu

Mahesh Nagarajan

Sauder School of Business, University of British Columbia, Vancouver, BC V6T 1Z2, Canada
email: mahesh.nagarajan@sauder.ubc.ca

In this paper, we describe the first computationally efficient policies for stochastic inventory models with lost sales and replenishment lead times that admit worst-case performance guarantees.

In particular, we introduce *dual-balancing policies* for lost-sales models that are conceptually similar to dual-balancing policies recently introduced for a broad class of inventory models in which demand is backlogged rather than lost. That is, in each period, we balance two opposing costs: the expected marginal holding costs against the expected marginal lost-sales cost. Specifically, we show that the dual-balancing policies for the lost-sales models provide a worst-case performance guarantee of 2 under relatively general demand structures. In particular, the guarantee holds for independent (not necessarily identically distributed) demands and for models with correlated demands such as the $AR(1)$ model and the multiplicative auto-regressive demand model. The policies and the worst-case guarantee extend to models with capacity constraints on the size of the order and stochastic lead times. Our analysis has several novel elements beyond the balancing ideas for backorder models.

Key words: Inventory, Approximation ; Dual-Balancing ; Algorithms; Lost Sales

MSC2000 Subject Classification: Primary: 90B05 , ; Secondary: 68W25 ,

OR/MS subject classification: Primary: inventory/production , approximation/heuristics ; Secondary: production/scheduling , approximation/heuristics

1. Introduction In this paper, we address one of the fundamental problems in stochastic inventory theory, the *single-item, single location, periodic-review, stochastic inventory control problem with lost sales*, which we refer to as the *lost-sales problem*. This problem has challenged researchers and practitioners for over five decades as very little is known about the structure of the optimal policy, and there are no known provably good heuristics even for the simplest settings. We build on ideas first proposed by Levi, Pál, Roundy and Shmoys [5]. They proposed what are called *dual-balancing policies* for a class of inventory models where unsatisfied demand is backlogged rather than lost. These policies have worst-case performance guarantees, that is, for each instance of the problem, the expected cost of the policy is guaranteed to be at most C times the optimal expected cost (for some constant C). In this paper, we discuss the implementation and the worst-case analysis of dual-balancing policies applied to inventory models with lost sales. These models have mathematical characteristics that are very different than the models in which excess demand is backlogged and thus require a fundamentally different and novel worst-case analysis. In particular, we shall describe the first computationally efficient policies for inventory models with lost sales that have a worst-case performance guarantee of 2. The analysis is based on several new ideas that we believe will contribute to future research in this domain.

Stochastic inventory theory provides streamlined models with the following common setting. The goal is to coordinate a sequence of orders over a planning horizon of finitely many discrete periods, aiming to supply a sequence of random demands over the planning horizon with minimum expected cost. The cost consists of a per-unit *ordering cost* for ordering supply units at the beginning of each period (with or without capacity constraints), a per-unit *holding cost* for carrying excess inventory from one period to the next, and a per-unit penalty cost for not satisfying demand on time. The dynamics of these models is as follows. At the beginning of each period, before the demand in this period is observed, a non-negative procurement order is placed with an outside supplier incurring a cost proportional to the number of units ordered. This order will arrive and become available after a *lead time* of several periods. The demand in that period is then observed and is satisfied to the maximum extent possible from the current *inventory*

on-hand. At the end of the period, two possible costs are incurred: excess units of inventory incur a proportional holding cost and unsatisfied units of demand incur a proportional penalty cost. The goal is to find an ordering policy that minimizes the overall expected costs over the entire horizon.

There are two common assumptions regarding unsatisfied demand at the end of a period. These different assumptions distinguish between two fundamentally different classes of models. In the first class of models, the assumption is that unsatisfied units of demand will stay in the system, each incurring a per-unit penalty cost for each period until it is satisfied. That is, unsatisfied demand is *backlogged* from period to period in a manner symmetric to excess inventory that is carried from period to period. These are called inventory models with *backlogged demands*. In the second class of models, which is the focus of this paper, each unsatisfied unit of demand is lost, i.e., it incurs a one-period penalty cost and then leaves the system. While these two classes of models are equivalent if the lead time is equal to zero, that is when the orders arrive instantaneously, they are fundamentally different for any positive lead time. In particular, the state-of-the-art knowledge on lost-sales models with lead times is very limited compared to the well understood models with backlogged demands.

Dynamic programming has been the most dominant paradigm in studying stochastic inventory models with lost sales and backlogged demand (see Zipkin [15] and Section 2.1 below for dynamic programming formulations of the lost-sales model). The optimization problem is defined recursively over time, using subproblems for each possible state of the system. In particular, the ordering decision in each period is made based on the available information at the beginning, which includes the joint conditional distribution of future demands, additional information that may be available by that period and the *pipeline vector*. The pipeline vector consists of the inventory on-hand at the beginning of the period and the quantities of the *outstanding orders* that were placed in past periods and have not yet arrived. Clearly, the pipeline vector has length equal to the lead time, which suggests that the state space of the corresponding dynamic program can grow exponentially fast with the length of the lead time.

However, it turns out that in models with backlogged demands, it is sufficient to consider only the *sum* of the inventory on-hand at the beginning of the period (or any backlogged demand) and the quantities of the outstanding orders. This sum is usually called the *inventory position* of the system. The intuition is that since all unsatisfied demands are backlogged, the impact of the decision made in the current period on the future costs depends only on the difference between the inventory position of the system (after ordering) and the cumulative demand over the lead time (to be realized). Moreover, the optimal policies in the models with backlogged demands have a simple form and are called *state-dependent base-stock policies*. In each period, there is a target inventory position level, referred to as the base-stock level, which is unaffected by the specific pipeline vector. If at the beginning of the period the inventory position is below the target base-stock level, we order up-to that target. If the inventory position at the beginning of the period is above the target base-stock level, no order is placed. The optimal base-stock levels can be computed by solving the corresponding dynamic program. Since the inventory position is a “sufficient statistic” for the pipeline vector, the computational complexity of dynamic programs for backorder models is insensitive to the lead time and is almost solely dictated by the complexity of the demand structure. We refer the reader to [15, 10, 5] for proofs of the optimality of base-stock policies and a discussion of the relevant literature regarding inventory models with backlogged demands.

In contrast, in systems with lost-sales systems and a positive lead time, the impact of the decision made in each period on the future costs is captured through a complicated mathematical expression that depends on the specific *sequence* of both the outstanding orders as well as the demands over subsequent periods. Specifically, the optimal decision in each period depends on the *entire* pipeline vector and not only on the inventory position as is the case in models with backlogged demands. As a result, the optimal policy in lost-sales models is significantly more complex and does not take the simple form of a base-stock policy and the inventory position is not a sufficient statistic for the pipeline vector. Moreover, the state space of the corresponding dynamic program grows exponentially fast with the lead time.

Due to the aforementioned difficulties, the literature on lost-sales models is limited. Karlin and Scarf [4] have been the first to study the optimal policies for models with lost sales and positive lead times. They have considered a lost-sales model with discrete finite and infinite horizon and with independent, identically distributed demands. They have shown that a base-stock policy cannot be optimal. Furthermore, for the case where the lead time is equal to one period, they have partially characterized the structure of the optimal policy. Specifically, they have shown that the optimal ordering quantity is a

decreasing function of the inventory on-hand at the beginning of the period, and is equal to zero outside a specified interval. Moreover, the rate of decrease (as a function of the inventory on-hand) is strictly smaller than one. With the additional assumption that demands are exponentially distributed, they have also presented a steady-state analysis of the dynamics of lost-sales systems that use base-stock policies. Morton [8] has extended the analysis of Karlin and Scarf to lost-sales models with deterministic lead times. He has shown that the optimal ordering policy is a function of the entire pipeline with the following characteristics: (a) there is a compact region around the origin (that is, all the components of the pipeline vector are zero) such that the order quantity is strictly positive if and only if the pipeline vector is in this region, (b) the order quantity decreases at a rate strictly between zero and one with respect to each component of the pipeline vector and (c) the rate of decrease in the order quantity per component is higher for components in the pipeline that are scheduled to arrive later in time. He has also derived upper and lower bounds on the probability that the optimal policy will have enough inventory to meet demand in the period in which the current order will arrive (a lead time ahead). Furthermore, he has used these bounds to derive upper and lower bounds on the optimal ordering policy as a function of the current pipeline vector. In a subsequent paper, Zipkin [17] has used state transformation techniques to establish simpler proofs for the structure of optimal policies in the lost-sales models discussed in this paper. Moreover, he has extended the results of Karlin and Scarf and Morton to models with capacity constraints on the size of the order, Markov modulated demands and stochastic lead times (with no order-crossing).

Morton [9] has considered *myopic policies* for lost-sales models, in which, in each period, an order is placed that minimizes the expected cost in the period in which this order arrives. There are other papers on lost-sales models like the ones by Nahmias [12] and Johansen [2] that propose different heuristics and present computational results on the performance of these heuristics. The computational experiments in all of these papers are restricted to instances where the lead time is short, equal to one or two periods or to models with extremely low demands. In a recent subsequent paper, Zipkin [16] presents computational experiments in which he tests the performance of several heuristics, including the dual-balancing policy described in this paper. He has focused on scenarios where the demands are independent and identically distributed; more specifically, they follow Poisson and Geometric distributions. Using state-reduction techniques, he is able to compute the optimal policy for instances with lead time equal to 4. Computing optimal policies with respect to instances with longer lead times seems very challenging. Moreover, to the best of our knowledge, there is no heuristic for lost-sales models that has been shown to perform well over a large bed of test problems of realistic size. Equally importantly, none of these papers provides a worst-case analysis of the proposed heuristics. Moreover, Levi, Pál, Roundy and Shmoys [5] have shown that the myopic policy for the lost-sales model even with lead time equal to zero does not have worst-case performance guarantees. Specifically, they have shown a class of examples for which the myopic policy is arbitrarily more expensive than the optimal policy. Reiman has considered a model with continuous time and with demand following a Poisson process and compared base-stock policies and policies that place an order in a fixed frequency [13].

In this paper we build on the recent work of Levi, Pál, Roundy and Shmoys [5] who have developed what are called dual-balancing policies for a class of uncapacitated stochastic inventory models with backlogged demands. These ideas have been extended to capacitated models [7] and multi-echelon models [6], again with backlogged demands. These dual-balancing policies are computationally efficient and have a worst-case performance guarantee of 2 for the respective models under general assumptions on the demand structure and the cost parameters. The dual-balancing policies are based on two novel ideas: a *marginal cost accounting approach* and *cost balancing* techniques. We note that the marginal cost accounting scheme is very different than the standard dynamic programming based cost accounting approach traditionally used to analyze these models. Using the marginal cost accounting approach, the dual-balancing policy is based on the repeated use of cost balancing techniques. In each period, two opposing (i.e. the holding and backlogging) costs are balanced. The worst-case analysis in the above three papers is heavily based on the mathematical properties of models with backlogged demands and uses a period-by-period amortization cost of the dual-balancing policy with the cost of the optimal policy. Crucial to the analysis is the fact that in backorder models, comparing the inventory positions of any two policies in a period provides sufficient information to analyze their respective performance a lead time ahead.

In this paper, we describe a dual-balancing policy for models with lost sales, which is conceptually similar to the dual-balancing policy for models with backlogged demands. However, the above-mentioned

analysis for models with backlogged demand is not applicable to models with lost sales. In particular, the inventory position *does not* provide sufficient information to compare the costs of different policies. In addition, a period-by-period amortization of the cost of the dual-balancing policy with the cost of the optimal policy does not seem useful. To overcome these difficulties, we describe a fundamentally different analysis which is based on two novel ideas. Rather than a period-by-period comparison, we use a *global* amortization of the lost-sales costs of the dual-balancing policy with the lost-sales costs of the optimal policy. In addition, we introduce a new concept called the *truncated inventory position* which generalizes the aforementioned notion of inventory position. As we have already mentioned, the inventory position in a certain period is defined to be the sum of the on-hand inventory at the beginning of the period plus *all* outstanding orders. The truncated inventory position is defined to be the sum of the on-hand inventory plus all the outstanding orders that have been ordered by a certain period, possibly *earlier* than the current period. In other words, the truncated inventory position accounts for the on-hand inventory and all the outstanding orders that will arrive by a certain period. The new concept of truncated inventory position is used to compare two policies in a lost-sales system. Our main result is that the dual-balancing policy for the lost-sales model has a worst-case performance guarantee of 2.

The worst-case analysis holds for models with relatively general demand structures. For example, it holds under the assumption that the demands in different periods are independent, not necessarily identically distributed (see Section 3 below for details). Moreover, the analysis also holds in many models in which the demands in different periods are correlated; specifically, it holds in the multiplicative autoregressive demand model and the $AR(1)$ model, which are commonly used in the literature. Finally, the policy and the worst-case analysis extends to models with stochastic lead times (under the assumption of “no-crossing of orders”) and to models with capacity constraints on the size of the order in each period.

We note that the dual-balancing policy can be computed efficiently in most if not all of the realistic scenarios. As an example, we focus attention on the case where the demands are independent integer-valued random variables with bounded support, and provide a detailed analysis of the running time of the dual-balancing policy. Dynamic programming approach seems to be computationally intractable, since the running time grows exponentially fast in the lead time. In contrast, we show that the dual-balancing policy can be computed in time polynomial in the number of periods and the length of the support of the demands.

The rest of the paper is organized as follows. In Section 2, we describe the lost-sales model and a dynamic programming formulation of the model. In Section 3, we describe a dual-balancing policy for lost-sales models and the new worst-case analysis under the assumption that the demands in different periods are independent. In Section 4, we discuss related computational issues of the dual-balancing policy. Finally, in Section 5 we describe several important extensions of the dual-balancing policy and the worst-case analysis to models with capacity constraints on the size of the order, models with stochastic lead times and to models with demand structures that allow correlation between demands in different periods.

2. The Lost-Sales Model In this section, we provide the mathematical formulation of the lost-sales model and introduce some of the notation used throughout the paper.

As a general convention, we distinguish between a random variable and its realization using capital letters and lower case letters, respectively. Script font is used to denote sets.

We consider a finite planning horizon of T periods numbered $t = 1, \dots, T$ (note that t and T are both deterministic). There is a sequence of stochastic demands that occur over the planning horizon, which are denoted by D_1, \dots, D_T , all of which have finite mean. We first assume that demands in different periods are independent of each other, though not necessarily identically distributed. In Section 5 we shall show that this assumption can be relaxed to allow several important structures of correlation between demands in different periods.

As part of the model, we assume that at the beginning of each period s , there is an observed *information set* that is denoted by f_s . The information set f_s contains all of the information that is available at the beginning of time period s . More specifically, the information set f_s consists of the realized demands (d_1, \dots, d_{s-1}) over the interval $[1, s)$ (in general f_s can contain additional information that became available by time period s). The information set f_s in period s is one specific realization in the set of all

possible realizations of the random vector $F_s = (D_1, \dots, D_{s-1})$. This set is denoted by \mathcal{F}_s . We consider only non-anticipatory policies, that is, in making a decision in period s , a feasible policy can use only the observed information set f_s .

In each period $s = 1, \dots, T$, a non-negative procurement order is placed from an outside supplier, incurring a per-unit ordering cost c_s . The order placed in period s will arrive and become available only after a positive *lead time*, denoted by L . We assume that L is a known positive integer, that is, an order placed in period s will arrive at the beginning of period $s + L$. (In Section 5, we will consider models where the lead times are stochastic.)

We now describe the dynamics of the lost-sales model. At the beginning of each period s , as a function of the observed information set f_s , we observe the joint (conditional) distribution of future demands (if demands in different periods are independent of each other, then the joint distribution is fixed regardless of the observed information set). At the beginning of period s , the system is characterized through the *pipeline vector*. The pipeline vector is denoted by \bar{P}_s and consists of L components. The L^{th} component is the *inventory on hand* (or on-hand inventory) available at the beginning of period s after the order placed L periods ago in period $s - L$ has arrived and before the demand in period s is realized. We denote the inventory on-hand at the beginning of period s by I_s . The other $L - 1$ components of the pipeline vector are the *outstanding orders* that have been placed in previous periods and have not yet arrived. Specifically, the j^{th} component of \bar{P}_s is equal to Q_{s-j} , the size of the order placed j periods ago, i.e., in period $s - j$ (for $j = 1, \dots, L - 1$). That is,

$$\bar{P}_s = (Q_{s-1}, \dots, Q_{s+1-L}, I_s).$$

Observe that at the beginning of period s all the components of the pipeline vector are known deterministically.

We next specify the sequence of events in each period s :

- (i) The order of q_{s-L} units placed in period $s - L$ arrives and the on-hand inventory is thus $i_s = (i_{s-1} - d_{s-1})^+ + q_{s-L}$. Observe that $(i_{s-1} - d_{s-1})^+$ is the inventory on-hand at the end of period $s - 1$.
- (ii) Following a given policy, q_s units are ordered ($q_s \geq 0$), and this incurs a cost of $c_s q_s$. Next the demand in period s is realized and is satisfied to the maximum extent possible from the inventory on-hand. Since unsatisfied demand is lost and leaves the system, the on-hand inventory decreases by $\min\{d_s, i_s\}$. In addition, we observe a new information set $f_{s+1} \in \mathcal{F}_{s+1}$.
- (iii) At the end of the period, costs are incurred. If $(i_s - d_s) > 0$ then we incur a total holding cost of $h_s(i_s - d_s)$ (this means that there is excess inventory that needs to be carried to time period $s + 1$). On the other hand, if $(i_s - d_s) < 0$ we incur a total lost-sales penalty cost of $p_s(d_s - i_s)$ (this means that in period s there is unsatisfied demand that is lost).

For ease of exposition, we first assume that the cost parameters are stationary, that is, for each $t = 1, \dots, T$, we have $h_t = h > 0$, $p_t = p > 0$ and $c_t = c \geq 0$. We further assume that $c = 0$. (The worst-case analysis presented below holds for any $c > 0$.) We will show that in fact the analysis allows us to have time-dependent holding costs parameters and non-increasing ordering and lost-sales penalty parameters. In particular, the analysis holds for models with stationary cost parameters and discount factor.

The goal is to find an ordering policy that minimizes the overall expected holding costs and lost-sales penalty costs over the entire horizon $[1, T]$. We consider only policies that are *non-anticipatory*, i.e., at time s , the information that a feasible policy can use consists only of f_s . Thus, for each feasible policy, given an information set f_s , the pipeline vector at time period s and the order quantity in period s are known deterministically.

2.1 Dynamic programming Formulation In this section, we discuss the dynamic programming formulation of the lost sales model and discuss the associated difficulties in the analysis.

Observe that in a lost-sales model the cost in period s depends on the inventory on-hand at the beginning of the period, that is, the expected cost in period s is equal to

$$E[h(I_s - D_s)^+ + p(D_s - I_s)^+].$$

Note that the decision made in period s of how many units to order affects only the costs over the interval $[s + L, T]$ (recall that the order placed in period s will arrive at the beginning of period $s + L$). Moreover, the impact of the decision in period s is captured through the effect it has on the inventory on-hand at the beginning of period $s + L$. Unfortunately, in lost-sales models, there is no tractable way to capture the impact of the decision in period s on the inventory on hand in period $s + L$. In particular, the impact of the decision made in period s on the inventory on-hand at the beginning of period $s + L$ depends on the specific *sequence* of both the outstanding orders at the beginning of period s and the realized demands over the interval $[s, s + L]$. Thus, the mathematical expressions of the dynamics of the lost-sales model are quite involved. As we have already seen, for each $t = 1, \dots, T$,

$$I_{t+1} = (I_t - D_t)^+ + Q_{t+1-L}, \quad (1)$$

which implies that the inventory on-hand in period $s + L$ depends on the decision of how many units to order in period s through a complicated recursive expression. Thus, the resulting dynamic program formulation depends on the *entire* observed pipeline vector. Let $V_s(\bar{p}_s, f_s) = V_s((q_{s-1}, \dots, q_{s+1-L}, i_s), f_s)$ be the optimal expected cost over the interval $[s + L, T]$ given an observed pipeline vector \bar{p}_s and an observed information set f_s . The recursion in the lost-sales model is

$$V_s((q_{s-1}, \dots, q_{s+1-L}, i_s), f_s) = \min_{q_s \geq 0} \{E[h(I_{s+L}(q_s) - D_{s+L})^+ + p(D_{s+L} - I_s(q_s))^+ | f_s] + E[V_{s+1}((q_s, q_{s-1}, \dots, q_{s+2-L}, (i_s - D_s)^+ + q_{s+1-L}), F_{s+1}) | f_s]\}, \quad (2)$$

where $I_{s+L}(q_s)$ is the inventory on-hand in period $s + L$ assuming that in period s , we have ordered q_s units. It is readily verified that the state space of the above dynamic program grows exponentially fast in the length of the lead time L even in simple cases where the demands in different periods are assumed to be independent and identically distributed. This implies that solving the above dynamic program is likely to be intractable except for cases with very small lead times. Moreover, the dynamic program does not provide much insight on the structure of the optimal policies and this a main reason why theoretical research on lost sales models is limited.

3. Dual-Balancing Policy for the Lost Sales Model In this section, we shall describe a *dual-balancing policy* for the lost-sales model, and then present a worst-case analysis that holds under relatively general assumptions on the demand distributions D_1, \dots, D_T . We shall show that under these assumptions, the dual-balancing policy has a worst-case performance guarantee of 2. That is, the expected cost of the policy is guaranteed to be at most twice the expected cost of an optimal policy. In this section, we describe the dual-balancing policy and its worst-case analysis in the case where demands in different periods are assumed to be independent of each other, though not necessarily identically distributed. In Section 5 we discuss several important extensions of the dual-balancing policy and the worst-case analysis to more general models.

3.1 Dual-Balancing Policy The policy for the lost-sales model is conceptually similar to the one proposed by Levi, Pál, Roundy and Shmoys for the model with backlogged demand [5]. That is, in each period s , conditioned on the observed information set f_s , we balance the (conditional) expected marginal holding cost incurred by the units ordered in that period over the entire horizon against the (conditional) expected lost-sales penalty cost incurred a lead time ahead in period $s + L$.

For a given policy P , let H_s^P be the marginal holding cost incurred by the units ordered in period s over the entire horizon, and let Π_s^P be the lost-sales penalty cost incurred in period $s + L$. The cost of policy P can then be expressed as

$$\mathcal{C}(P) = \sum_{s=1}^{T-L} (H_s^P + \Pi_s^P),$$

ignoring the marginal holding cost of units ordered before period 1 and the lost-sales penalty costs over the interval $[1, L]$ that are identical for every policy. However, the expressions of H_s^P and Π_s^P are different in the lost-sales model, and are significantly more complex compared to the corresponding expressions in the models with backlogged demands. Recall that I_t^P is the on-hand inventory in period t after the order placed in period $t - L$ has arrived and before the demand in period t has occurred. We have already seen that, for each $t = 1, \dots, T - 1$,

$$I_{t+1}^P = (I_t^P - D_t)^+ + Q_{t+1-L}^P, \quad (3)$$

(where Q_j^P for $j \leq 0$ are given as an input). Observe that $(I_t^P - D_t)^+$ is the inventory on-hand at the end of period t and Q_{t+1-L}^P is the order arriving at the beginning of period $t + 1$. Assuming without loss of generality that supply units are consumed on a first-ordered-first-consumed basis, we conclude that the Q_s^P units ordered in period s will be consumed only after all the $(I_{s+L-1} - D_{s+L-1})^+$ units that were on-hand at the beginning of period $s + L$ (just before the order placed in period s has arrived) are consumed. This leads to the following expression for the marginal holding cost incurred by the Q_s^P units ordered in period s :

$$H_s^P = \sum_{t=s+L}^T h(Q_s^P - (D_{[s+L,t]} - (I_{s+L-1}^P - D_{s+L-1})^+)^+). \quad (4)$$

Similarly, we express

$$\Pi_s^P = p(D_{s+L} - I_{s+L}^P)^+ = p(D_{s+L} - (Q_s^P + (I_{s+L-1}^P - D_{s+L-1})^+)^+), \quad (5)$$

where the second equality follows from Equation (3) above. Equations (4) and (5) can be easily adapted to capture time-dependent cost parameters. In addition, we can incorporate a linear ordering cost $c_s Q_s^P$ into Equation (4) above.

For each $s = 1, \dots, T - L$ and an observed information set $f_s \in F_s$, define the functions $l_s^B(q_s^B) = E[H_s^B(q_s^B)|f_s]$ and $\pi_s^B(q_s^B) = E[\Pi_s^B(q_s^B)|f_s]$. As in the dual-balancing policy for the model with backlogged demand [5], in each period s , conditioned on the observed information set f_s , we order $q_s^B = q'_s$ to balance $l_s^B(q'_s) = E[H_s^B(q'_s)|f_s] = \pi_s^B(q'_s) = E[\Pi_s^B(q'_s)|f_s]$. It is readily verified that, conditioned on f_s and the resulting pipeline vector \bar{p}_s^B , the functions l_s^B and π_s^B depend only on q_s^B . Moreover, l_s^B is an increasing (convex) function of q_s^B , which is equal to 0 if $q_s^B = 0$ and goes to infinity as q_s^B goes to infinity. In addition, π_s^B is a (convex) decreasing function of q_s^B , which admits a non-negative value for $q_s^B = 0$ and is going to 0 as q_s^B goes to infinity. If fractional orders are allowed the function l_s^B and π_s^B are continuous and thus q'_s is well defined. (Later we shall discuss the case where orders are restricted to be integral, and demands are integer-valued random variables.)

The intuition behind the idea of repeatedly balancing the functions π_s^B and l_s^B above is that in the lost-sales model there are two underlying opposing risks, the risk of *under ordering* incurring lost-sales penalty cost and the risk of *over ordering* incurring holding costs. Balancing these two risks seems to be very effective and computationally attractive. Surprisingly, this idea works significantly better than minimizing the sum of the two functions. We also note that the dual-balancing policy can be implemented in an on-line manner. That is, the decision made in each period is not affected by any future decision of the policy, but only by the currently observed information set. This seems like an essential property if one wishes to avoid the burden of solving huge dynamic programs. However, unlike *myopic policy*, which in each period aims to minimize only the expected cost a lead time ahead, the dual-balancing policy does look ahead make use of available information about the future evolution of the system.

Integral orders and integer-valued demands. Next we discuss the case where the demands are integer-valued random variables and the order quantity in each period is restricted to be an integer. We briefly describe a *randomized dual-balancing policy* using ideas identical to ones used in [5, 7].

In this case, the functions $l_s^B(q_s^B)$ and $\pi_s^B(q_s^B)$ are initially defined only for integer values. Their piecewise linear interpolations preserve the monotonicity (and convexity) properties discussed in Section 3. The problem is that the balancer q'_s is likely to be fractional. Instead we consider the two consecutive integers $q_s^1 \leq q'_s \leq q_s^2$. It is clear that $q'_s = \lambda q_s^1 + (1 - \lambda)q_s^2$ for some $0 < \lambda < 1$. We now order q_s^1 with probability λ and q_s^2 with probability $1 - \lambda$.

3.2 Analysis - Independent Demands Given the dual-balancing policy for the lost-sales model, we define Z_t to be the random balanced cost in period t , i.e., $Z_t = E[H_t^B|\mathcal{F}_t] = E[\Pi_t^B|\mathcal{F}_t]$ (for each $t = 1, \dots, T - L$). Using an identical proof to the one in [5], we obtain the following lemma.

LEMMA 3.1 *The expected cost of the dual-balancing policy is equal to twice the sum of expectations of the Z_t variables, i.e., $E[C(B)] = 2 \sum_{t=1}^{T-L} E[Z_t]$.*

The worst-case analysis of the dual-balancing policy in models with backlogged demand [5] is based on a period by period amortization of the cost of the dual-balancing policy against the optimal policy.

This is done by comparing the respective inventory positions of the the two policies, in each period [5]. In contrast, it is well-known [15] that looking only on the inventory position is not sufficient to make optimal decisions in lost-sales models.

Similarly, unlike the analysis of the models with backlogged demands, comparing the inventory positions of the dual-balancing policy and *OPT* in period s does not seem to provide ‘sufficient’ information about period $s + L$. For example, consider a lost-sales model with $L = 1$, where in period t the pipeline vector of policy 1 is $\bar{p}_t^1 = (3, 10)$ (i.e., on-hand inventory equal to 10 and an order of 3 units placed in period t) and the pipeline vector of policy 2 is $\bar{p}_t^2 = (4, 1)$ (i.e., on-hand inventory equal to 1 and an order of 4 units placed in period t). In period t the inventory position of policy 1 is $y_t^1 = 13$, higher than the inventory position of policy 2, which is $y_t^2 = 5$. However, if the demand in period t is greater than 10, then policy 2 has greater on-hand inventory in period $t + 1$ (4 units) than policy 1, which is left only with 3 units on-hand. Conversely, if the demand in period t is no greater than 9, then policy 1 has on-hand inventory in period $t + 1$ no smaller than that of policy 2.

The above example suggests that the period-by-period amortization scheme of the cost of the dual-balancing policy against the cost of *OPT*, based on the inventory position as used in the backlogging analysis, does not seem to be useful when applied to the lost-sales model. (In models with backlogged demand if one policy has a higher inventory position in period s , it will have higher on-hand inventory a lead time ahead in period $s + L$.) To overcome this difficulty, the analysis presented below incorporates two novel ideas.

We use a *global* amortization of costs, that is, we compare the *overall* cost of the dual-balancing policy to that of *OPT*, where the comparison is not necessarily period-by-period. In addition, we introduce the new concept of *truncated inventory position*, which is defined as follows. For each period $s = 1, \dots, T$, the truncated inventory position with respect to period t (where $t \in [s - L, s]$), is defined to be the sum of the inventory on-hand in period s plus all outstanding orders *placed by time period t* . Let Y_{st} denote the truncated inventory position in period s with respect to period t , that is,

$$Y_{st} = I_s + \sum_{j=s+1-L}^t Q_j. \quad (6)$$

Observe that the truncated inventory position Y_{st} refers to the sum of the on-hand inventory in period s plus all outstanding orders that *will arrive by time period $t + L$* . Note that we consider a period t earlier than s which implies that all the orders that arrive by time period $t + L$ are already known at time period s . Specifically, $Y_{ss} = Y_s$ is the traditional inventory position defined earlier in Section 2, and $Y_{s,s-L} = I_s$ is the on-hand inventory at the beginning of period s . The truncated inventory position is a generalization of the traditional inventory position concept commonly used in inventory theory (see Figure 1). Due to the complex mathematical structure of lost-sales models, the effect of the decision made in the current period on future costs is very hard to quantify. The truncated inventory position provides a more tractable way to analyze this effect; specifically, the effect of the current ordering decision on the on-hand inventory a lead time ahead. Moreover, it turns out that the concept of the truncated inventory position provides the ‘right’ framework for comparing between the pipeline vectors of any two policies; specifically, *OPT* and the dual-balancing policy. Thus, a central part of the worst-case analysis presented below is based on this new concept. We believe that it will have more applications in other settings.

The worst-case analysis in the model with backlogged demand is based on comparing the (traditional) inventory position of the dual-balancing policy and *OPT* in each period t , i.e., comparing Y_{tt}^B and Y_{tt}^{OPT} . Instead, in the lost-sales model, the analysis will be based on comparing the respective truncated inventory positions Y_{st}^B and Y_{st}^{OPT} in each period $s \in [t, t + L]$. That is, in each period $s \in [t, t + L]$, we compare the respective number of units already ordered by the dual-balancing policy and *OPT* that will be available by time period $t + L$ (see Figure.2).

Let \mathcal{T}_H be the set of all periods $t \leq T - L$ such that the truncated inventory of the dual-balancing policy with respect to period t is strictly smaller than the respective truncated inventory position of




Figure 1: The truncated inventory position in period s ($L=2$).

OPT , for each period $s \in [t, t + L]$. That is,

$$\mathcal{T}_H = \{t \leq T - L : \forall s \in [t, t + L], Y_{st}^B < Y_{st}^{OPT}\}. \quad (7)$$

Let \mathcal{T}_Π be the complement of \mathcal{T}_H , i.e., the set of periods $t \leq T - L$ for which there exists some $s \in [t, t + L]$ where the truncated inventory position of the dual-balancing policy with respect to period $t + L$ is no smaller than the respective truncated inventory position of OPT . That is,

$$\mathcal{T}_\Pi = \{t \leq T - L : \exists s \in [t, t + L] \text{ with } Y_{st}^B \geq Y_{st}^{OPT}\}. \quad (8)$$

Recall that in the lost-sales model, having a higher inventory position in period t *does not* guarantee higher on-hand inventory in period $t + L$. Moreover, for certain realizations of the demands over the interval $[t, t + L]$, the truncated inventory position of the dual-balancing policy with respect to period t might be higher than the respective truncated inventory position of OPT in some periods and lower in others. In fact, it is possible to observe an *alternating behavior*, where the relation between the truncated inventory position of the dual-balancing policy and that of OPT may change several times over the interval. More precisely, for some period t and $s \in [t, t + L]$, we will say that the respective truncated inventory position of the dual-balancing policy with respect to period t and that of OPT alternate in period s if one of the following events occur

$$[Y_{st}^B < Y_{st}^{OPT}] \cap [Y_{s+1,t}^B \geq Y_{s+1,t}^{OPT}],$$

or

$$[Y_{st}^B \geq Y_{st}^{OPT}] \cap [Y_{s+1,t}^B < Y_{s+1,t}^{OPT}].$$

That is, in the two consecutive periods s and $s + 1$, the inequalities relating the truncated inventory positions of the dual-balancing policy and that of OPT alternate.

For each $t \in \mathcal{T}_H$, we know that in each period over the interval $[t, t + L]$, OPT had (strictly) more units available by time period $t + L$. In particular, there is no alternation in the respective relation between the truncated inventory position of the dual-balancing policy with respect to period t and that of OPT over the interval $[t, t + L]$. On the other hand, for each $t \in \mathcal{T}_\Pi$, there was at least one period over that interval when the dual-balancing policy had units available by time period $t + L$ at least as many as OPT had. Note that this does not necessarily imply alternations (e.g., when the truncated inventory position of the dual-balancing policy with respect to period t is higher in period t and throughout $(t, t + L]$), nor does it exclude more than one alternation (i.e., it is possible that the respective truncated inventory position of OPT and that of the dual-balancing policy will alternate several times over $[t, t + L]$).

Next we state and prove two key lemmas that will show how to amortize the cost of the dual-balancing policy against the cost of OPT . The corresponding two lemmas hold with probability 1, i.e., for each sample path of the demands D_1, \dots, D_T or equivalently, for each $f_T \in \mathcal{F}_T$. (In the statements and proofs




Figure 2: Evolution of the truncated inventory position with respect to period t over $[t, t + L]$ ($L = 3$)

of these lemmas we shall omit the expression ‘with probability 1’.) In the first of these lemmas we will show that the overall holding cost incurred by OPT , denoted by H^{OPT} is greater than the holding costs incurred by units ordered by the dual-balancing policy in periods $t \in \mathcal{T}_H$.

LEMMA 3.2 *The holding cost incurred by OPT is greater than the holding cost incurred in the dual-balancing policy by units ordered in periods $t \in \mathcal{T}_H$, i.e., $H^{OPT} \geq \sum_{t \in \mathcal{T}_H} H_t^B$.*

PROOF. Recall that by definition $Y_{t+L,t}^P = I_{t+L}^P$. However, this implies that, for each $t \in \mathcal{T}_H$, we have $I_{t+L}^B < I_{t+L}^{OPT}$, i.e., the on-hand inventory of OPT in period $t+L$ is higher than that of the dual-balancing policy. We have already seen that the on-hand inventory at the beginning of period $t+L$, just before the units ordered in period t have arrived, is equal to $(I_{t+L-1}^P - D_{t+L-1})^+$. In particular,

$$I_{t+L}^B = Q_t^B + (I_{t+L-1}^B - D_{t+L-1})^+ < I_{t+L}^{OPT}.$$

Without loss of generality, we assume that supply units are consumed on a first-ordered-first-consumed basis. We can then associate an index to each unit of supply currently on-hand according to the number of units on-hand to be consumed prior to that unit (where units ordered in the same period are sorted arbitrarily). Note that since we allow fractional orders, the supply units are defined infinitesimally. In particular, the Q_t^B units ordered by the dual-balancing policy in period t are indexed in period $t+L$ in the range

$$((I_{t+L-1}^B - D_{t+L-1})^+, (I_{t+L-1}^B - D_{t+L-1})^+ + Q_t^B]. \quad (9)$$

Since $t \in \mathcal{T}_H$ and the on-hand inventory of OPT in period $t+L$ is higher, we conclude that in period $t+L$ there exist supply units on-hand in OPT with the same range of indices as in (9). We now *match* pairs of units of supply with the same respective index (in period $t+L$) in the dual-balancing policy and OPT , respectively. In particular, in period $t+L$ we match the supply units that are indexed in the above range in OPT to the Q_t^B units ordered by the dual-balancing policy in period t (see also Figure 3).

Observe that until the I_{t+L}^B units on-hand at the beginning of period $t+L$ will be consumed, neither the dual-balancing policy nor OPT incur lost-sales costs. Moreover, since the demands over $[t+L, T]$ are the same for OPT and the dual-balancing policy, it is clear that each pair of respective matched supply units of OPT and the dual-balancing policy will incur the same holding cost over $[t+L, T]$, for each sample path of demands D_{t+L}, \dots, D_T . Since each pair of units are consumed at the same time period, it is readily verified that each supply unit of OPT can be matched to at most one supply unit of the dual-balancing policy. This concludes the proof. \square

Note that the above proof still holds for time-dependent holding cost parameters and positive non-increasing per-unit ordering cost parameters, where the per-unit ordering cost is incorporated into the marginal expected holding cost and is balanced against the marginal expected lost-sales penalty cost.

In the second lemma, we amortize the lost-sales penalty costs of the dual-balancing policy which are associated with periods $t \in \mathcal{T}_\Pi$. In the proof of this lemma, we use a global amortization rather than a period-by-period one. For each $t \in \mathcal{T}_\Pi$, we know that there exists some period $s \in [t, t+L]$ such that the truncated inventory position of the dual-balancing policy with respect to period t is no smaller than the one of OPT , i.e., $Y_{st}^B \geq Y_{st}^{OPT}$. However, as we have already observed, this does not guarantee that in period $t+L$ the inventory on-hand of the dual-balancing policy is no smaller than the one of OPT . That is, it is still possible to have $I_{t+L}^B < I_{t+L}^{OPT}$, which implies that we can not amortize the lost-sales penalty cost incurred by the dual-balancing policy in period $t+L$ against the respective cost of OPT in this period. The next lemma shows that in this case, period $t+L$ belongs to an *interval of periods* over which the lost-sales penalty costs incurred by OPT are higher than the respective lost-sales penalty costs incurred by the dual-balancing policy. This leads to a global amortization of the cost of the dual-balancing policy with the cost of OPT .

LEMMA 3.3 *The lost-sales penalty incurred by OPT , denoted by Π^{OPT} , is greater than the lost-sales penalty costs of the dual-balancing policy which are associated with periods $t \in \mathcal{T}_\Pi$, i.e., $\Pi^{OPT} \geq \sum_{t \in \mathcal{T}_\Pi} \Pi_t^B$.*




Figure 3: Matched supply units in period $t + L$ where $t \in \mathcal{T}_H$.

PROOF. Consider the following random partition of the periods $L + 1, \dots, T$. For each realization of demands d_1, \dots, d_T , consider the resulting realization of the set \mathcal{T}_Π , and partition the periods in the following way. Start in period T and look for the latest period $t \in \mathcal{T}_\Pi$ with the property that period $t + L$ is not marked (initially all periods are unmarked) and $i_{t+L}^B < i_{t+L}^{OPT}$ (we abuse the notation and use \mathcal{T}_Π to denote the deterministic set of periods resulting from the realized demands). If no such t exists then we terminate. If such a period exists, let t' be that period and let $w_{t'}$ be the earliest period in $[t', t' + L]$ for which the truncated inventory position of the dual-balancing policy with respect to t' is no smaller than the respective truncated inventory position of OPT . That is, $w_{t'} = \min\{j \in [t', t' + L] : y_{jt'}^B \geq y_{jt'}^{OPT}\}$ (observe that $w_{t'}$ is the realization of a random variable, denoted by $W_{t'}$, which is defined for each period $t' \in \mathcal{T}_\Pi$). By our assumption t' does belong to \mathcal{T}_Π , hence $w_{t'}$ is indeed well-defined. We now *mark* all the periods in $[w_{t'}, t' + L]$. Next we continue recursively over the periods $1, \dots, w_{t'} - 1$. That is, we look for the latest $t \leq w_{t'} - 1$ such that $t \in \mathcal{T}_\Pi$ and with the property that $t + L$ is unmarked and $i_{t+L}^B < i_{t+L}^{OPT}$ and repeat the above.

The above procedure induces a random partition of the periods $L + 1, \dots, T$ into marked and unmarked periods, respectively. Let \mathcal{M} be the (random) set of all marked periods. In particular, this random partition induces a partition of the set \mathcal{T}_Π into periods $s \in \mathcal{T}_\Pi$ such that $s + L \in \mathcal{M}$, i.e., $s + L$ is marked and periods $s \in \mathcal{T}_\Pi$ such that $s + L$ is not marked. First consider the latter set. For each period $s \in \mathcal{T}_\Pi$ such that $s + L \notin \mathcal{M}$, we know that $I_{s+L}^B \geq I_{s+L}^{OPT}$, for if not $s + L$ would have been marked. This implies that for all periods $\{s \in \mathcal{T}_\Pi : s + L \notin \mathcal{M}\}$, we have $\Pi_s^B \leq \Pi_s^{OPT}$.

Now consider all the periods $\{s \in \mathcal{T}_\Pi : s + L \in \mathcal{M}\}$. Since all marked intervals are disjoint, it is sufficient to show that, for each marked interval of the type $[W_{t'}, t' + L]$, the overall lost-sales penalty costs incurred by OPT over that interval are higher than the respective lost-sales penalty costs incurred by the dual-balancing policy over that interval. In particular, this will imply that the lost-sales costs of the dual-balancing policy associated with periods in the set $\{s \in \mathcal{T}_\Pi : s + L \notin \mathcal{M}\}$ are lower than the

lost-sales penalty costs incurred by OPT in periods which do not belong to \mathcal{M} , and that the lost-sales penalty costs of the dual-balancing policy associated with periods in the set $\{s \in \mathcal{T}_{\Pi} : s + L \in \mathcal{M}\}$ are smaller than the lost-sales costs incurred by OPT in periods that belong to \mathcal{M} (see Figure 3.4). The proof of the lemma will then follow.

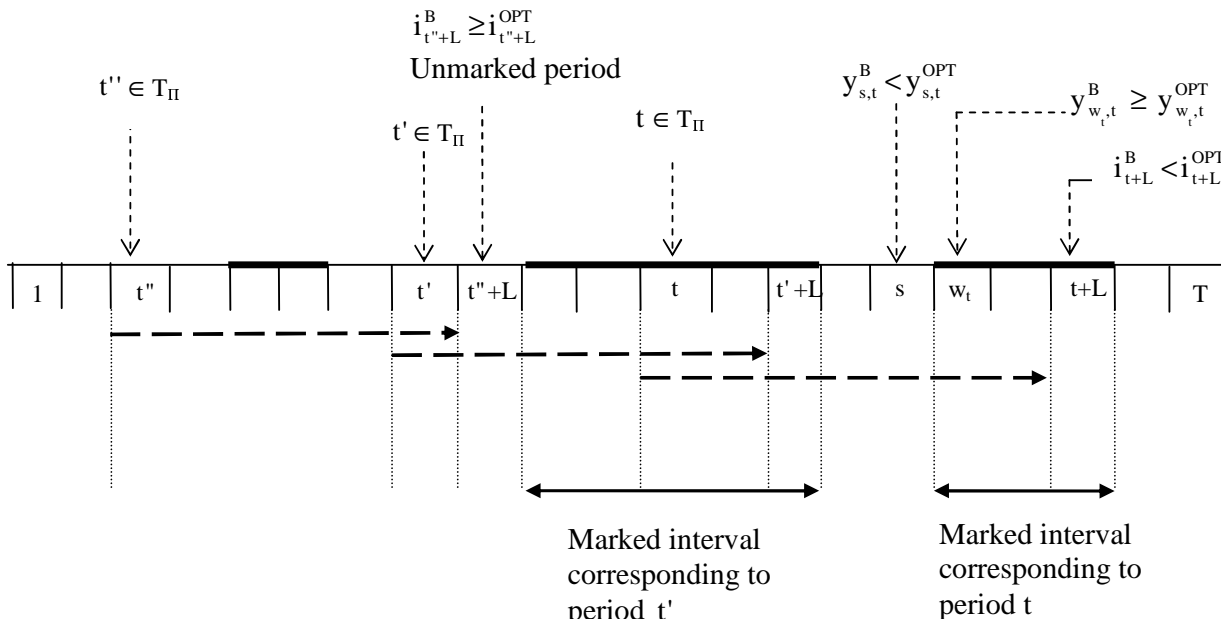


Figure 4: The partition into marked and unmarked periods. Dashed arrow denotes the lead time interval starting at a certain period. Solid arrow denote a marked interval. For each $t \in \mathcal{T}_{\Pi}$ with $t + L$ unmarked (e.g., period t''), we have $i_{t+L}^B \geq i_{t+L}^{OPT}$.

For each such interval $[W_{t'}, t' + L]$, we know that $Y_{W_{t'}, t'+L}^B \geq Y_{W_{t'}, t'+L}^{OPT}$, and $I_{t'+L}^B < I_{t'+L}^{OPT}$. Moreover, the difference between the lost-sales penalty cost incurred by the dual-balancing policy in period $t' + L$ and the respective lost-sales penalty cost incurred by OPT in that period, is bounded by p times the difference between the respective on-hand inventory levels. That is,

$$\Pi_{t'}^B - \Pi_{t'}^{OPT} \leq p(I_{t'+L}^{OPT} - I_{t'+L}^B). \tag{10}$$

Next we use the following identity that is valid for every feasible policy. For each t, s such that $t \leq s \leq t + L$,

$$I_{t+L} = Y_{st} - D_{[s, t+L]} + \frac{\Pi_{[s-L, t]}}{p}, \tag{11}$$

where $\Pi_{[s-L, t]}$ is the cumulative lost-sales penalty costs over the interval $[s, t + L]$, i.e., $\Pi_{[s-L, t]} = \sum_{j=s-L}^{t-1} \Pi_j$. Equation (11) describes the dynamics of a model with lost sales. Specifically, the on-hand inventory in period $t + L$ is equal to the truncated inventory position in time period s with respect to period t minus the cumulative demand over the interval $[s, t + L]$ plus the cumulative lost sales over that interval. Observe that $D_{[s, t+L]} - \frac{\Pi_{[s-L, t]}}{p}$ is the number of supply units consumed by the demand over the interval $[s, t + L]$. Now consider Equation (11) for periods t' and $W_{t'}$ applied to OPT and the dual-balancing policy, respectively, and substitute into Equation (10). We get that

$$\begin{aligned} \Pi_{t'}^B - \Pi_{t'}^{OPT} &\leq \Pi_{[W_{t'}-L, t']}^{OPT} - \Pi_{[W_{t'}, t']}^B + p(Y_{W_{t'}, t'}^{OPT} - Y_{W_{t'}, t'}^B) \\ &\leq \Pi_{[W_{t'}-L, t']}^{OPT} - \Pi_{[W_{t'}-L, t']}^B. \end{aligned} \tag{12}$$

The last inequality follows from the fact that $Y_{W_{t'}, t'}^B \geq Y_{W_{t'}, t'}^{OPT}$. We now get that

$$\begin{aligned} \sum_{t \in \mathcal{T}_\Pi} \Pi_t^B &= \sum_{t: t \in \mathcal{T}_\Pi, t+L \in \mathcal{M}} \Pi_t^B + \sum_{t: t \in \mathcal{T}_\Pi, t+L \notin \mathcal{M}} \Pi_t^B \\ &\leq \sum_{t: t \in \mathcal{T}_\Pi, t+L \in \mathcal{M}} \Pi_t^{OPT} + \sum_{t: t \in \mathcal{T}_\Pi, t+L \notin \mathcal{M}} \Pi_t^{OPT} \leq \Pi^{OPT}. \end{aligned} \quad (13)$$

This concludes the proof of the lemma. \square

We note that Lemma 3.3 holds also in the case where there are time-dependent lost-sales penalty parameters p_1, \dots, p_T , as long as they are non-increasing. The proof is almost identical, but now $w_{t'}$ is defined to be the latest period $j \in [t', t' + L]$, such that the cumulative lost sales of the dual-balancing over $[j, t' + L]$ is no higher than the corresponding lost sales of OPT over that interval. (The proof of Lemma 3.3 implies that the newly defined $w_{t'}$ does exist.) This enables us to amortize the lost sales incurred by the dual-balancing policy, in each period t such that $t - L \in \mathcal{T}_\Pi$, with lost sales incurred by OPT in periods earlier than t . (Specifically, for each period $s \in [w_{t'}, t' + L]$, we amortize the lost sales of the dual-balancing in period s with lost sales of OPT incurred in periods $[w_{t'}, s]$.) In particular, the lemma is valid in models with discounted costs.

Lemmas 3.2 and 3.3 imply that

$$H^{OPT} + \Pi^{OPT} \geq \sum_{t \in \mathcal{T}_H} H_t^B + \sum_{t \in \mathcal{T}_\Pi} \Pi_t^B.$$

Taking expectation we get that

$$E[\mathcal{C}(OPT)] \geq E\left[\sum_t (\mathbb{1}(t \in \mathcal{T}_H) \cdot H_t^B + \mathbb{1}(t \in \mathcal{T}_\Pi) \cdot \Pi_t^B)\right]. \quad (14)$$

However, as we have already seen, in the lost-sales model the truncated inventory positions of the dual-balancing policy and OPT with respect to period t can alternate over the interval $[t, t + L)$ from higher to lower. Thus, unlike the analysis of model with backlogged demand [5], conditioning on some $f_t \in \mathcal{F}_t$ does not necessarily realize the indicators $\mathbb{1}(t \in \mathcal{T}_H)$ and $\mathbb{1}(t \in \mathcal{T}_\Pi)$ above. That is, it is possible that in period t we still do not know whether $t \in \mathcal{T}_H$ or $t \in \mathcal{T}_\Pi$.

Instead, we will condition on the events $[t \in \mathcal{T}_H]$ and $[t \in \mathcal{T}_\Pi]$, respectively, and get that

$$\begin{aligned} E[\mathcal{C}(OPT)] &\geq \sum_t E[E[\mathbb{1}(t \in \mathcal{T}_H) \cdot H_t^B | \mathcal{F}_t] + E[\mathbb{1}(t \in \mathcal{T}_\Pi) \cdot \Pi_t^B | \mathcal{F}_t]] \\ &= \sum_t E[Pr(t \in \mathcal{T}_H | \mathcal{F}_t) \cdot E[H_t^B | (\mathcal{F}_t, t \in \mathcal{T}_H)] + Pr(t \in \mathcal{T}_\Pi | \mathcal{F}_t) \cdot E[\Pi_t^B | (\mathcal{F}_t, t \in \mathcal{T}_\Pi)]]. \end{aligned} \quad (15)$$

However, by conditioning on $[t \in \mathcal{T}_H]$ and $[t \in \mathcal{T}_\Pi]$, respectively, we consider information that supersedes the original information set $f_t \in \mathcal{F}_t$ based on which the dual-balancing policy has made the ordering decision at the beginning of period t . That is, $E[H_t^B | (\mathcal{F}_t, t \in \mathcal{T}_H)]$ and $E[\Pi_t^B | (\mathcal{F}_t, t \in \mathcal{T}_\Pi)]$ might not be equal to $E[H_t^B | \mathcal{F}_t] = E[\Pi_t^B | \mathcal{F}_t] = Z_t$. In particular, the problem arises for information sets $f_t \in \mathcal{F}_t$ for which $y_{tt}^B < y_{tt}^{OPT}$ and the conditional probabilities (conditioning on f_t) that $[t \in \mathcal{T}_H]$ and $[t \in \mathcal{T}_\Pi]$ are both positive (If this is not the case, then we know whether $t \in \mathcal{T}_H$ or $t \in \mathcal{T}_\Pi$ already at the beginning of period t while observing f_t .)

Next we will show that if the demands in different periods are independent of each other, then the two inequalities

$$Pr(t \in \mathcal{T}_H | \mathcal{F}_t) \cdot E[H_t^B | (\mathcal{F}_t, t \in \mathcal{T}_H)] \geq Pr(t \in \mathcal{T}_H | \mathcal{F}_t) \cdot E[H_t^B | \mathcal{F}_t] \quad (16)$$

and

$$Pr(t \in \mathcal{T}_\Pi | \mathcal{F}_t) \cdot E[\Pi_t^B | (\mathcal{F}_t, t \in \mathcal{T}_\Pi)] \geq Pr(t \in \mathcal{T}_\Pi | \mathcal{F}_t) \cdot E[\Pi_t^B | \mathcal{F}_t] \quad (17)$$

hold with probability 1, and this together with Equation (15) and the fact that $Z_t = E[H_t^B | \mathcal{F}_t] = E[\Pi_t^B | \mathcal{F}_t]$ imply that the dual-balancing has a worst-case performance guarantee of 2. Equations (16) and (17) imply that conditioning also on the events $[t \in \mathcal{T}_H]$ and $[t \in \mathcal{T}_\Pi]$, respectively, implies that

the respective expected costs are even higher than what was expected at period t conditioning only on f_t . Note that this is the only part of the analysis that requires additional assumptions on the demand distributions (beyond having finite mean). In Section 5, we shall generalize the analysis, and show that Inequalities (16) and (17) hold under several other demand structures that incorporate correlation between demands in different periods. Intuitively, we require that the demands do not have a certain ‘bad’ property. That is, we would like to exclude a situation where high demands over a certain interval of periods, say (j', j) , imply low demands over the rest of the horizon $[j, T]$. Indeed, if the demands are independent, this ‘bad’ situation is excluded.

For each period $t = 1, \dots, T - L$ and $s \in [t, t + L)$, let A_{st} be the event that at the beginning of period s the truncated inventory position of OPT with respect to period t is higher than the one of the dual-balancing policy, while at the beginning of period $s + 1$ the truncated inventory position of the dual-balancing policy with respect to period t is no smaller than the one of OPT . That is,

$$A_{st} = [Y_{st}^B < Y_{st}^{OPT}] \cap [Y_{s+1,t}^B \geq Y_{s+1,t}^{OPT}].$$

Observe that conditioning on an information set f_t such that $y_{tt}^B < y_{tt}^{OPT}$, then $t \in \mathcal{T}_\Pi$ only if the event A_{jt} occurs for some $j \in [t, t + L)$. In the next lemma we characterize some of the properties of the event A_{st} defined above.

LEMMA 3.4 *For each period $t = 1, \dots, T - L$ and $s \in [t, t + L)$, let A_{st} be as defined above. Suppose that the event A_{st} occurred. Then,*

(i) *The cumulative amount of orders placed by the dual-balancing policy over the interval $[s + 1 - L, t]$ is higher than the corresponding amount of orders of OPT over that interval, i.e., $\sum_{j=s+1-L}^t Q_j^B \geq \sum_{j=s+1-L}^t Q_j^{OPT}$.*

(ii) *The inventory on-hand of OPT at the beginning of period s exceeds that of the dual-balancing policy by more than $\Delta Q_s = \sum_{j=s+1-L}^t Q_j^B - \sum_{j=s+1-L}^t Q_j^{OPT}$, i.e., $I_s^{OPT} > I_s^B + \Delta Q_s$.*

(iii) *The event A_{st} can be expressed as*

$$[Y_{st}^B < Y_{st}^{OPT}] \cap [\Delta Q_s \geq 0] \cap [D_s > I_s^{OPT} - \Delta Q_s].$$

(iv) *The dual-balancing policy has incurred positive lost sales in period s , and hence its on-hand inventory at the beginning of period $s + 1$ is equal to the size of the order placed in period $s + 1 - L$, denoted by Q_{s+1-L}^B . That is, $A_{st} \subseteq [\Pi_{s-L}^B > 0] \subseteq [I_{s+1}^B = Q_{s+1-L}^B]$.*

PROOF. Recall Equation (6) that, for each policy P , we have $Y_{st}^P = I_s^P + \sum_{j=s+1-L}^t Q_j^P$. Assume that (i) does not hold, i.e., that A_{st} has occurred and that $\sum_{j=s+1-L}^t Q_j^B < \sum_{j=s+1-L}^t Q_j^{OPT}$. Since $Y_{st}^B < Y_{st}^{OPT}$, we conclude that $I_s^B - I_s^{OPT} < \sum_{j=s+1-L}^t Q_j^{OPT} - \sum_{j=s+1-L}^t Q_j^B$. However, it is readily verified that this implies that the inventory on-hand of the dual-balancing policy at the end of period s , $(I_s^B - D_s)^+$, does not exceed the respective inventory on-hand of OPT , $(I_s^{OPT} - D_s)^+$, by more than $\sum_{j=s+1-L}^t Q_j^{OPT} - \sum_{j=s+1-L}^t Q_j^B$. That is,

$$(I_s^B - D_s)^+ - (I_s^{OPT} - D_s)^+ < \sum_{j=s+1-L}^t Q_j^{OPT} - \sum_{j=s+1-L}^t Q_j^B,$$

which implies that $Y_{s+1,t}^B < Y_{s+1,t}^{OPT}$ and leads to a contradiction. The proof of (i) then follows.

The proof of (ii) follows from (i) and the fact that $Y_{st}^B < Y_{st}^{OPT}$.

It is now clear that given (i) and (ii) above, the event A_{st} is equivalent to the event $[D_s > I_s^{OPT} - \Delta Q_s]$, which implies (iii). Finally, (ii) and (iii) imply (iv). \square

In the next two lemmas we show that if the demands are independent of each other, then the Inequalities (16) and (17) do hold. (We again omit the statement ‘with probability 1’ as long as the context is clear.)

LEMMA 3.5 *Assume that D_1, \dots, D_T are independent of each other. Then for each period $t = 1, \dots, T - L$, we have $\Pr(t \in \mathcal{T}_H | \mathcal{F}_t) \cdot E[H_t^B | \mathcal{F}_t] \leq \Pr(t \in \mathcal{T}_H | \mathcal{F}_t) \cdot E[H_t^B | (\mathcal{F}_t, t \in \mathcal{T}_H)]$.*

PROOF. Consider some information set $f_t \in \mathcal{F}_t$. If $Pr(t \in \mathcal{T}_H | f_t)$ is equal either to 0 or to 1, then there is nothing to prove. Suppose now that $0 < Pr(t \in \mathcal{T}_H | f_t) < 1$. In particular, f_t is such that $y_{tt}^B < y_{tt}^{OPT}$. It is now sufficient to show that $E[H_t^B | f_t] \leq E[H_t^B | (f_t, t \in \mathcal{T}_H)]$.

Let W_t be the earliest period $s \in [t, t+L]$ such that $Y_{st}^B \geq Y_{st}^{OPT}$ if such period exists and $W_t = -1$ otherwise. It is readily verified that, conditioning on f_t , the event $[W_t = -1]$ is equivalent to the event $[t \in \mathcal{T}_H]$. Note that the specific information set f_t being considered implies that $Pr(W_t = t | f_t) = 0$. Using conditional expectations, we now write

$$E[H_t^B | f_t] = Pr(t \in \mathcal{T}_H | f_t) \cdot E[H_t^B | (f_t, t \in \mathcal{T}_H)] + \sum_{j=t+1}^{t+L} Pr(W_t = j | f_t) \cdot E[H_t^B | (f_t, W_t = j)]. \quad (18)$$

It is now sufficient to show that, for each $j = t+1, \dots, t+L$, $E[H_t^B | f_t] \geq E[H_t^B | (f_t, W_t = j)]$. This together with Equation (18) imply that indeed $E[H_t^B | f_t] \leq E[H_t^B | (f_t, t \in \mathcal{T}_H)]$ from which the proof follows.

Recall, that conditioning on f_t we know already the size of the order placed by the dual-balancing policy at the beginning of period t . Let $\phi_{jt}(i_j)$ be the expected marginal holding cost of the $q_t^B = q_t'$ units ordered by the dual-balancing policy in period t conditioning on f_t and on the inventory on-hand at the beginning of period j equal to i_j (for each $j = t+1, \dots, t+L$). That is, $\phi_{jt}(i_j) = E[H_t^B | (f_t, I_j^B = i_j)]$. Observe that $E[H_t^B | f_t] = E[\phi_{jt}(I_j^B) | f_t]$. Since the demands in different periods are independent of each other, it is readily verified that D_j, \dots, D_T are independent of I_j^B and \mathcal{F}_t . It follows from Equations (3) and (4) that ϕ_{jt} is increasing in i_j . Moreover, $I_j^B | f_t \geq q_{j-L}^B$ (where q_{j-L}^B is the size of the order arriving at the beginning of period j , which conditioning of f_t is known deterministically). Thus, $\phi_{jt}(I_j^B) | f_t \geq \phi_{jt}(q_{j-L}^B)$ and $E[\phi_{jt}(I_j^B) | f_t] \geq \phi_{jt}(q_{j-L}^B)$. Finally, observe that the event $f_t \cap [W_t = j]$ is also independent of the demands D_j, \dots, D_T . In addition, Lemma 3.4 above implies that the event $f_t \cap [W_t = j]$ is contained in the event $A_{j-1,t}$, i.e., contained in the event $[I_j^B = Q_{j-L}]$. We conclude that $E[H_t^B | (f_t, W_t = j)] = \phi_{jt}(q_{j-L}^B)$. This concludes the proof of the lemma. \square

LEMMA 3.6 *Assume D_1, \dots, D_T are independent of each other. Then for each period $t = 1, \dots, T-L$, we have $Pr(t \in \mathcal{T}_\Pi | \mathcal{F}_t) \cdot E[\Pi_t^B | \mathcal{F}_t] \leq Pr(t \in \mathcal{T}_\Pi | \mathcal{F}_t) \cdot E[\Pi_t^B | (f_t, t \in \mathcal{T}_\Pi)]$.*

PROOF. Consider an information set $f_t \in \mathcal{F}_t$. If $Pr(t \in \mathcal{T}_\Pi | f_t) = 0$ or $Pr(t \in \mathcal{T}_\Pi | f_t) = 1$, then there is nothing to prove. Suppose that $0 < Pr(t \in \mathcal{T}_\Pi | f_t) < 1$. In particular, f_t is such that $y_{tt}^B < y_{tt}^{OPT}$. Let W_t be the same random variable as defined in Lemma 3.5 above. We express

$$\begin{aligned} E[\Pi_t^B | (f_t, t \in \mathcal{T}_\Pi)] &= \sum_{j=t}^{t+L} Pr(W_t = j | f_t, t \in \mathcal{T}_\Pi) \cdot E[\Pi_t^B | (f_t, W_t = j, \mathcal{T}_\Pi)] \\ &= \sum_{j=t}^{t+L} Pr(W_t = j | f_t, t \in \mathcal{T}_\Pi) \cdot E[\Pi_t^B | (f_t, W_t = j)]. \end{aligned}$$

Since $Pr(W_t = t | f_t) = 0$, it is sufficient to show, for each $j = t+1, \dots, t+L$, that $E[\Pi_t^B | f_t] \leq E[\Pi_t^B | (f_t, W_t = j)]$.

For each $j = t+1, \dots, t+L$, let $\psi_{jt}(i_j)$ be the expected lost-sales cost incurred by the dual-balancing policy in period $t+L$ conditioning on f_t and on the inventory on-hand at the beginning of period j equal to i_j . That is, $\psi_{jt}(i_j) = E[\Pi_t^B | (f_t, I_j^B = i_j)]$. Similar to Lemma 3.5 above, we conclude that ψ_{jt} is decreasing in i_j . This implies that

$$E[\Pi_t^B | f_t] = E[\psi_{jt}(I_j^B) | f_t] \leq \psi_{jt}(q_{j-L}^B).$$

However, we have already observed that the event $f_t \cap [W_t = j]$ is independent of the demands D_j, \dots, D_{t+L} , and is contained in the event $A_{j-1,t}$. From Lemma 3.4, we conclude that $E[\Pi_t^B | (f_t, W_t = j)] = \psi_{jt}(q_{j-L}^B)$, from which the proof of the lemma follows. \square

Equation (15) and Lemmas 3.5 and 3.6 imply the following theorem.

THEOREM 3.1 *Consider the lost-sales model with independent demands, time-dependent holding cost parameters and non-increasing ordering and lost-sales penalty parameters. Then the dual-balancing policy has a worst-case performance guarantee of 2.*

Randomized dual-balancing policy . Next we extend the worst-case analysis of the randomized dual-balancing policy in the case where order quantities are restricted to be integers and the demands are integer-valued random variables. Observe that now at the beginning of period s (conditioning on the observed information set f_s) we know q_s^1 and q_s^2 above deterministically. However, the actual size of the order is still random. Using the same definition of Z_t above, it is readily verified that Lemma 3.1 is valid and that $E[\mathcal{C}(B)] = 2 \sum_t E[Z_t]$. Next we modify the definition of the sets \mathcal{T}_H and \mathcal{T}_Π and define them with respect to the truncated inventory position of the dual-balancing policy assuming that it orders q_t^2 units (though it might end up ordering only q_t^1). Denote this truncated inventory by \tilde{Y}_{st}^B i.e., this is the truncated inventory position of the dual-balancing policy in period s assuming that it has ordered q_t^2 units in period t . In particular, $0 \leq \tilde{Y}_{st}^B - Y_{st}^B \leq 1$. Specifically, we define

$$\mathcal{T}_H = \{t \leq T - L : \forall s \in [t, t + L] \tilde{Y}_{st}^B \leq Y_{st}^{OPT}\},$$

and

$$\mathcal{T}_\Pi = \{t \leq T - L : \exists s \in [t, t + L] \tilde{Y}_{st}^B > Y_{st}^{OPT}\}.$$

It is readily verified, that for each $t \in \mathcal{T}_H$, we have $I_{t+L}^B \leq I_{t+L}^{OPT}$, and that, for each $t \in \mathcal{T}_\Pi$ there exists $s \in [t, t + L]$ such that $Y_{st}^B \geq Y_{st}^{OPT}$. Thus, all the arguments used in the proofs of Lemmas 3.2 and 3.3 are still valid. The proofs of Lemmas 3.4, 3.5 and 3.6 directly apply.

4. Computational Issues Next we discuss several computational issues regarding the implementation of the dual-balancing policy. The main goal of this discussion is to highlight the fundamental difference in the computational efforts required by the dual-balancing policy compared to the traditional dynamic programming approach. The formal notion of computational efficiency in the context of stochastic optimization models is, by far, less clear than in deterministic optimization models. In particular, in stochastic optimization models, one can make different assumptions regarding the way the probability distributions are specified, and the oracles that are available, and each set of assumptions leads to a different analysis. Thus, a detailed analysis can be done only in the context of a concrete scenario. In what follows, we shall restrict the discussion to two related issues:

- (i) What oracles are needed in general for the implementation of the dual-balancing policy, and how this is compared to the dynamic programming approach and the myopic policy.
- (ii) A detailed analysis of the computational effort required to implement the dual-balancing policy in the important special case, where the demands D_1, \dots, D_T are independent (not necessarily identically distributed) integer-valued random variables with support within $\{0, \dots, M\}$, where $M \in \mathbb{N}$. The goal is to highlight the fact that, in concrete and important scenarios, the dual-balancing policy can be implemented efficiently, whereas the dynamic programming approach does not seem to be tractable.

To implement the dual-balancing policy, we need to compute the *balancer* q'_s in each period $s = 1, \dots, T - L$. As a result, the running time of the dual-balancing policy is of the order T times the complexity of computing q'_s in each period. The fact that q'_s lies in the intersection of a decreasing $(\pi_s^B(q_s))$ and an increasing function $(l_s^B(q_s))$ suggests that bisection search methods will be very efficient for computing q'_s as long as the functions l_s^B and π_s^B above can be evaluated efficiently.

From Equations (4) and (5) above we can see that in order to evaluate the functions l_s^B and π_s^B it is sufficient to be able to evaluate the distributions of $D_{[s+L,t]}$ (for each $t = s + L, \dots, T$) and $(I_{s+L-1}^B - D_{s+L-1})^+$. The latter distribution is usually the most challenging one to evaluate. However, in period s conditioning on the observed information set f_s , we already know deterministically the pipeline vector \bar{p}_s^B including i_s^B , the on-hand inventory at the beginning of period s . Thus, we can use Equation (3) to recursively compute the distributions $I_{s+1}^B, I_{s+2}^B, \dots, I_{s+L-1}^B$. Specifically, if there are efficient oracles to evaluate the cumulative demand distributions $D_{[s+L,t]}$, for each $s = 1, \dots, T - L$ and $t \geq s + L$, and the distributions $(X - D_s)^+$, for each $s = 1, \dots, T$ and a random variable X , then the functions π_s^B and l_s^B can be evaluated efficiently.

This suggests that, in most if not all of the common scenarios, there exist efficient ways to evaluate the functions l_s^B and π_s^B . Note that the computational effort required grows moderately as the lead time L grows. This is in contrast to the traditional dynamic programming approach that is very sensitive to the lead time since the corresponding state space grows exponentially fast in the lead time L , even in

the simplest scenarios (see Section 2.1 above). In fact, the exponential growth of the state space makes it very hard to compute the optimal policies in models with lead times longer than 4 [16].

Moreover, the requirement to have an oracle that evaluates the distributions $(X - D_s)^+$ above is essential even for computing a simple myopic policy. As we have already mentioned, this policy aims, in each period s , to minimize the overall expected cost in period $s + L$ [15, 16]. However, for computing the myopic policy one needs to evaluate the distribution $(I_{s+L-1}^B - D_{s+L-1})^+$.

Next we discuss the important special case in which the demands D_1, \dots, D_T are independent (not necessarily identically distributed) integer-valued random variables with support over $\{0, \dots, M\}$, where $M \in \mathbb{N}$. We shall show how to compute the dual-balancing policy in time polynomial in T , L and M .

First, observe that like any reasonable policy the dual-balancing policy will never order more than M units in a period. This implies that the inventory on-hand of the dual-balancing policy will never exceed $M(L + 1)$. Moreover, the functions π_s^B and l_s^B are piecewise linear with break points in the integers $0, \dots, M$. Thus, if there is an efficient way to evaluate the functions π_s^B and l_s^B , one can compute the balancer q'_s using $O(\log M)$ calls for these functions. We simply apply a bisection search on $\{0, \dots, M\}$ to find the two consecutive integers q_s^1 and $q_s^2 = q_s^1 + 1$ such that $\pi_s^B(q_s^1) \geq l_s^B(q_s^1)$ and $\pi_s^B(q_s^2) \leq l_s^B(q_s^2)$. (Since the functions are piecewise linear, the balancer q'_s can be expressed as a convex combination of q_s^1 and q_s^2 , and then the randomized dual-balancing can be applied.) It is now sufficient to show that one can evaluate efficiently the functions π_s^B and l_s^B .

Focus on period s . Suppose that the cumulative distributions $D_{[s+L,t]}$ and the distribution of $(I_{s+L-1}^B - D_{s+L-1})^+$ are given explicitly, i.e., by specifying the values in the support and the corresponding probabilities. Since D_{s+L}, \dots, D_T have support within $\{0, \dots, M\}$ it follows that the support of $D_{[s+L,t]}$ is within $\{0, \dots, TM\}$. Since all the orders and demands are integral, assuming that i_0 is also an integer implies that $(I_{s+L-1}^B - D_{s+L-1})^+$ is an integer-valued random variables with support within $\{0, 1, \dots, M(L+1)\}$. From Equations (4) and (5) it follows that, for a given order quantity q_s^B , computing the functions $\pi_s^B(q_s^B)$ and $l_s^B(q_s^B)$ can be done in $O(T^2LM^2)$ time. Specifically, l_s^B consists of a sum of $O(T)$ elements, one for each period $j = s + L, \dots, T$. Each element j in the sum can be computed by enumerating over all the possible $M(L+1)$ values of $(I_{s+L-1}^B - D_{s+L-1})^+$ and the possible TM values of $D_{[s+L,j]}$. For each j , this can be done in $O(TLM^2)$ time. Thus, the decision, in each period s , can be computed in $O(T^2LM^2 \log M)$ time.

Next we show how to construct the explicit distributions $D_{[s+L,t]}$ and $(I_{s+L-1}^B - D_{s+L-1})^+$. However, if the distribution $D_{[s+L,j]}$ is given explicitly, we can construct the distribution $D_{[s+L,j+1]}$ in $O(TM^2)$ time. Thus, the cumulative demand distributions can be computed recursively in $O(T^2M^2)$ time. Recall that conditioning on f_s we know i_s^B . The recursive Equation (3) and the fact that D_1, \dots, D_T are integer-valued random variables with support within $\{0, \dots, M\}$ imply that, for each $j = 0, \dots, L-1$, the random variable $(I_{s+j}^B - D_{s+j})^+$ has at most $(j+1)M + 1$ values in its support. Moreover, if $(I_{s+j}^B - D_{s+j})^+$ is given explicitly, we can compute the explicit distribution of $(I_{s+j+1}^B - D_{s+j+1})^+$ in $O(jM^2)$. Since we use the recursion of Equation (3) $L-1$ times, it follows that the distribution of $(I_{s+L-1}^B - D_{s+L-1})^+$ can be expressed explicitly in $O(M^2L^2)$. This implies the following theorem.

THEOREM 4.1 *Consider the lost-sales model with integer-valued independent demands D_1, \dots, D_T with support within $\{0, \dots, M\}$, for some $M \in \mathbb{N}$. Then the dual-balancing policy can be computed in $O(T^2LM^2 \log M)$ time.*

We note that computing an optimal policy in the above scenarios seems extremely hard unless the lead time is very short. There are many other scenarios in which the dual-balancing policy can be computed efficiently. The specific analysis requires detailed specifications of the underlying assumptions. We believe that in most of the common scenarios, the dual-balancing policy proposed above will be straightforward to implement both computationally and conceptually.

5. Extensions

5.1 Infinite Horizon In this subsection, we briefly discuss the implementation of the dual-balancing policy in models with infinite horizon. Conceptually, the dual-balancing policy can still be implemented. However, the computational effort involved depends on the maximum number of periods that a unit

ordered in some period can stay in inventory before consumed. (This will affect the computation of the marginal holding costs.) Finally we note that the same worst-case analysis described in Section 3 above still holds in infinite horizon models with average or discounted cost.

5.2 Capacitated Model In this subsection, we shall describe how to extend the policy and the worst-case analysis to the case where there exists a capacity constraint on the size of the order in each period. Specifically, we assume that, for each period $t = 1, \dots, T$, there is a given nonnegative capacity $u_t \geq 0$, and that the size of the order placed in that period can not exceed the capacity, i.e., $q_t^P \leq u_t$.

We first describe a marginal cost accounting scheme for lost-sales models with capacity constraints on the size of the order in each period. This scheme follows closely the ideas introduced by Levi, Roundy, Shmoys and Truong in the context of capacitated models with backlogged demand [7]. The idea is to associate with the decision of how many units to order in each period s , all the future lost-sales costs that are *forced* by this decision. Specifically, let \bar{q}_s^P be the slack capacity in period s following some policy P , i.e., $\bar{q}_s^P = u_s - q_s^P$ is the unused capacity in period s (for $s = 1, \dots, T - L$). For each $s = 1, \dots, T - L$ and $t \geq s$, let $\tilde{\Pi}_{st}^P$ denote the overall lost-sales costs incurred in period t by a policy that follows policy P until period s and then orders up to capacity in each of the periods $s + 1, \dots, T - L$ (i.e., it orders up to u_j for each period $j = s + 1, \dots, T - L$). Period s is associated with all the lost-sales costs incurred in period t that could have been avoided by ordering up to capacity in period s (instead of ordering q_s^P). That is, the lost-sales cost in period t associated with period s is equal to the difference $\tilde{\Pi}_{st}^P - \tilde{\Pi}_{s-1,t}^P$. This is referred to as the *forced lost-sales cost in period t due to the decision made in period s* . In addition, let $\tilde{\Pi}_{0t}^P$ (for each $t = L + 1, \dots, T - L$) be the lost-sales costs incurred in period t if the orders in all the periods are up to capacity. It is readily verified that $\tilde{\Pi}_{0t}^P$ captures costs that are unavoidable by every policy. Next define $\tilde{\Pi}_s^P$ to be the overall lost-sales costs associated with period s (for each $s = 1, \dots, T - L$):

$$\tilde{\Pi}_s^P = \sum_{t=s+L}^T (\tilde{\Pi}_{st}^P - \tilde{\Pi}_{s-1,t}^P). \tag{19}$$

We will sometimes refer to the difference $\tilde{\Pi}_{st}^P - \tilde{\Pi}_{s-1,t}^P$ as the *contribution* of period t to the lost-sales costs associated with period s . It is readily verified that $\tilde{\Pi}_s^P$ is at most $p\bar{q}_s$, that is, period s can not be associated with more than \bar{q}_s lost units of demand, since we could have ordered at most \bar{q}_s additional units in that period. By ordering in period s up to capacity, we can avoid part of these lost-sales costs up to $p\bar{q}_s^B$.

Observe that, for each $t = L + 1, \dots, T$, we have

$$\sum_{s=1}^{t-L} (\tilde{\Pi}_{st}^P - \tilde{\Pi}_{s-1,t}^P) + \tilde{\Pi}_{0t}^P = \tilde{\Pi}_{t-L,t}^P = \Pi_{t-L}^P. \tag{20}$$

Equation (20) above implies that the overall contributions of period t to the lost-sales costs associated with periods $1, \dots, t - L$ is equal to the lost-sales cost incurred in period t minus the term $\tilde{\Pi}_{0t}^P$ which is again the lost sales incurred in period t by any feasible policy. Thus, the above lost-sales cost accounting scheme is valid and exact (see Figure 5.5). In particular, the cost of each feasible policy P can be written as

$$\mathcal{C}(P) = \sum_{s=1}^{T-L} (H_s^P + \tilde{\Pi}_s^P), \tag{21}$$

ignoring the holding costs incurred by units ordered prior to period 1, lost-sales costs over $[1, L]$ and the term $\sum_{t=L+1}^T \tilde{\Pi}_{0t}^P$ which are common to every feasible policy.

Using the above marginal cost accounting scheme we construct a dual-balancing policy which is conceptually identical to the policy in the uncapacitated case. That is, in each period s , conditioned on the observed information set $f_s \in \mathcal{F}_s$, we order $q_s^B = q'_s$ to balance the functions $l_s^B(q_s^B) = E[H_s^B(q_s^B)|f_s]$ and $\tilde{\pi}_s^B(q_s^B) = E[\tilde{\Pi}_s^B(q_s^B)|f_s]$, i.e., we balance

$$l_s^B(q'_s) = E[H_s^B(q'_s)|f_s] = \tilde{\pi}_s^B(q'_s) = E[\tilde{\Pi}_s^B(q'_s)|f_s].$$

t	d_t	q_t	\bar{q}_t	$i_t = (i_{t-1} - d_{t-1})^+ + q_{t-1}$	Lost sales in period t	$\tilde{\Pi}_t$	Π_t
1	2	3	5	5	0	4	2
2	8	5	3	6	2 2→1	3	0
3	3	4	2	5	0	2	4
4	13	-	-	6	7 2→1 3→2 2→3	0	0

Figure 5: A specific sample-path of demands and orders ($L = 1$, $p = 1$). An arrow indicates an allocation of units of shortage to a period. For example, in period 4 there are 7 units of lost sales, out of which 2 units of shortage are associated with period 3 (i.e., $\tilde{\Pi}_{34} - \tilde{\Pi}_{24} = 2$), 3 units are associated with period 2 (i.e., $\tilde{\Pi}_{24} - \tilde{\Pi}_{14} = 3$) and 2 units are associated with period 1 (i.e., $\tilde{\Pi}_{14} - \tilde{\Pi}_{04} = 2$). The last two columns compare the lost-sales costs associated with each period according to the traditional and marginal cost accounting, respectively. In particular, we again balance the expected marginal holding costs incurred by the units ordered in period s against the expected marginal future lost-sales costs forced by the decision in period s . Note that the function $\tilde{\pi}_s^B$ is nonnegative at $q_s^B = 0$ and is equal to 0 at $q_s^B = u_s$. Thus, the balancer q_s^B above is again well-defined. Note that the function $\tilde{\pi}_s^B$ is more complicated to evaluate compared to π_s^B , the period expected lost-sales cost incurred in period $s + L$, defined above for the uncapacitated model. However, in most common scenarios it is still relatively straightforward to evaluate it accurately and efficiently. Observe that in the uncapacitated model, all the capacities are equal to ∞ and thus, $\bar{q}_s^P = \infty$. It is now readily verified that in this case $\tilde{\Pi}_s^P = \Pi_s^P$. That is, the marginal lost-sales cost accounting scheme for the capacitated model is in fact a generalization of the traditional lost-sales cost accounting scheme for uncapacitated models.

The worst-case analysis is very similar to the analysis presented in Section 3 above for the uncapacitated model. In particular, defining $Z_t = E[H_t^B | \mathcal{F}_t] = E[\tilde{\Pi}_t^B | \mathcal{F}_t]$, it is readily verified that Lemma 3.1 is still valid. That is, $E[\mathcal{C}(B)] = 2 \sum_t E[Z_t]$. We also keep the definitions of the sets \mathcal{T}_H and \mathcal{T}_Π the same as before. It is easy to see that Lemma 3.2 above is still valid, that is, $H^{OPT} \geq \sum_{s \in \mathcal{T}_H} H_s^B$. Next we state and prove a lemma analogous to Lemma 3.3.

LEMMA 5.1 *The lost-sales penalty incurred by OPT, denoted by Π^{OPT} , is greater than the lost-sales penalty costs of the dual-balancing policy which are associated with periods $s \in \mathcal{T}_\Pi$, i.e.,*

$$\Pi^{OPT} \geq \sum_{s \in \mathcal{T}_\Pi} \tilde{\Pi}_s^B.$$

PROOF. As in Lemma 3.3 we first describe a random partition of the periods $L+1, \dots, T$. Specifically, consider a realization of the demands d_1, \dots, d_T . Look for the latest period t with $i_t^B < i_t^{OPT}$ that contributes to the lost-sales costs associated with some period $s \in \mathcal{T}_\Pi$. That is, $\tilde{\Pi}_{st}^B - \tilde{\Pi}_{s-1,t}^B > 0$ for some $s \in \mathcal{T}_\Pi$. If no such t exists terminate. Otherwise, let t' be that period and let s' be the latest period in \mathcal{T}_Π within the interval $[1, t' - L]$. Note that since t' contributes lost-sales costs only to periods within $[1, t' - L]$ including some period $s \in \mathcal{T}_\Pi$ within that interval, s' above is well-defined. However, it is not necessary that t' contributes lost-sales costs to s' itself. Let $w_{s'}$ be the earliest period in $[s', s' + L]$, such that $Y_{w_{s'}, s'}^B \geq Y_{w_{s'}, s'}^{OPT}$ (since s' is in \mathcal{T}_Π , $w_{s'}$ is well-defined). Next mark all the periods in $[w_{s'}, t']$. As in Lemma 3.3 we repeat the above for the interval $[1, w_{s'}]$.

This induces a partition of the periods that contributes to the lost-sales costs associated with periods $s \in \mathcal{T}_\Pi$ into marked and unmarked periods. Consider first an unmarked period t that contributes to the lost-sales costs associated with periods $s \in \mathcal{T}_\Pi$. By the definition of the partition described above, it is readily verified that $I_t^B \geq I_t^{OPT}$, thus, $\Pi_{t-L}^{OPT} \geq \Pi_{t-L}^B$ (i.e., the lost sales incurred by OPT in period t is higher than the respective lost-sales costs of the dual-balancing policy). Moreover, Equation (20) implies that the lost-sales cost incurred by OPT in period t is higher than all the lost-sales costs incurred by the dual-balancing policy in period t , which are associated with periods $1, \dots, t - L$. In particular, the lost-sales cost OPT incurs in period t is higher than all of the contributions of period t to periods $s \in \mathcal{T}_\Pi$. It is now sufficient to show that the lost-sales costs incurred by OPT over an interval $[w_{s'}, t']$ are higher than the lost-sales costs incurred by the dual-balancing policy over that interval, which are associated

with periods $s \in \mathcal{T}_\Pi$. Observe that by the definition of the partition all of these periods are within the interval $[1, s']$. For each such interval, we have (we again abuse notation and use s' , t' and $w_{s'}$ though they are random variables),

$$\begin{aligned} \frac{\Pi_{t'-L}^{OPT}}{p} &= (D_{t'} - I_{t'}^{OPT})^+ \geq D_{t'} - I_{t'}^{OPT} \\ &= D_{t'} - (Y_{w_{s'}, s'}^{OPT} + \sum_{j \in (s', t'-L]} q_j^{OPT} - D_{[w_{s'}, t']}) + \frac{\Pi_{[w_{s'}-L, t'-L]}^{OPT}}{p}, \end{aligned}$$

where the equality follows again from the identity $I_t^{OPT} = Y_{w_s, s}^P + \sum_{j \in (s, t-L]} q_j^{OPT} - D_{[w_s, t]} + \frac{\Pi_{[w_s-L, t-L]}^{OPT}}{p}$ for each t , $s \leq t - L$ and $w_s \in [s, s + L]$. Taking $\frac{\Pi_{[w_s-L, t-L]}^{OPT}}{p}$ to the left-hand side, we get

$$\begin{aligned} \frac{\Pi_{[w_{s'}-L, t'-L]}^{OPT}}{p} &\geq D_{t'} - (Y_{w_{s'}, s'}^B + \sum_{j \in (s', t'-L]} u_j - D_{[w_{s'}, t']}) \\ &= D_{t'} - (Y_{w_{s'}, s'}^B + \sum_{j \in (s', t'-L]} u_j - D_{[w_{s'}, t']}) + \frac{\Pi_{[w_{s'}-L, s']}^B}{p} + \sum_{r \in (s'+L, t')} \frac{\tilde{\Pi}_{s'r}^B}{p} \\ &+ \frac{\Pi_{[w_{s'}-L, s']}^B}{p} + \sum_{r \in (s'+L, t')} \frac{\tilde{\Pi}_{s'r}^B}{p} = \frac{\Pi_{[w_{s'}-L, s']}^B}{p} + \sum_{r \in (s'+L, t')} \frac{\tilde{\Pi}_{s'r}^B}{p}. \end{aligned}$$

The first inequality follows from the fact that $Y_{w_{s'}, s'}^B \geq Y_{w_{s'}, s'}^{OPT}$ and $u_j \geq Q_j^{OPT}$. In the first equality we just add and subtract $\frac{\Pi_{[w_{s'}-L, s']}^B}{p} + \sum_{r \in (s'+L, t')} \frac{\tilde{\Pi}_{s'r}^B}{p}$. Here $\frac{\Pi_{[w_{s'}-L, s']}^B}{p}$ is as before the lost sales incurred by the dual-balancing policy over the interval $[w_{s'}, s' + L]$ and $\sum_{r \in (s'+L, t')} \frac{\tilde{\Pi}_{s'r}^B}{p}$ is the lost sales incurred over $(s' + L, t')$ by a policy that follows the dual-balancing policy until period s' and then orders up to capacity in periods $s' + 1, \dots, T - L$. The last equality follows from the fact that

$$D_{t'} - (Y_{w_{s'}, s'}^B + \sum_{j \in (s', t'-L]} u_j - D_{[w_{s'}, t']}) + \frac{\Pi_{[w_{s'}-L, s']}^B}{p} + \sum_{r \in (s'+L, t')} \frac{\tilde{\Pi}_{s'r}^B}{p}$$

is equal to the lost sales that the above policy incurs in period t' , i.e., equal to $\tilde{\Pi}_{s't'}^B$. In particular, observe that this policy has the same truncated inventory position in period $w_{s'}$ with respect to period s' , and over the time interval $[w_{s'}, s' + L]$ it incurs exactly $\frac{\Pi_{[w_{s'}-L, s']}^B}{p}$ units of lost sales (this policy follows the dual-balancing until period s').

Finally, we have already seen that the part of the lost-sales costs incurred by the dual-balancing policy over the time interval $[w_{s'}, s' + L]$ that is associated (by the marginal cost accounting scheme) with ordering decisions in other periods can not exceed the actual lost-sales cost incurred over this interval, i.e., $\Pi_{[w_{s'}-L, s']}^B$ (see also Equation (20) above). Moreover, from Equation (19) it follows that the part of the lost-sales costs incurred by the dual-balancing policy in some period $t \in [s' + L, t']$ that is associated (by the marginal cost accounting scheme) with the an ordering decision in some period $s \in [1, s']$ can not exceed $\tilde{\Pi}_{st}^B$, which is the lost sales incurred in period t by a policy that follows the dual-balancing policy until period s and then orders up to capacity in each of the subsequent periods. However, by definition this is at most $\tilde{\Pi}_{s't}^B$, which is similarly, the lost-sales cost incurred in period t by a policy that follows the dual-balancing policy until period s' , and only then starts to order up to capacity. The last step of the proof is identical to Equation (13) in the proof of Lemma 3.3 above (see Figure 5.6). □

Similar to the uncapacitated model discussed in Section 3, the proof of the Lemma 5.1 can be extended to models with non-increasing lost-sales penalty parameters. Observe that Lemmas 3.4 and 3.5 still hold. In addition, it can be verified that Lemma 3.6 still holds with respect to $\tilde{\Pi}_t^B$ (instead of Π_t^B), i.e.,

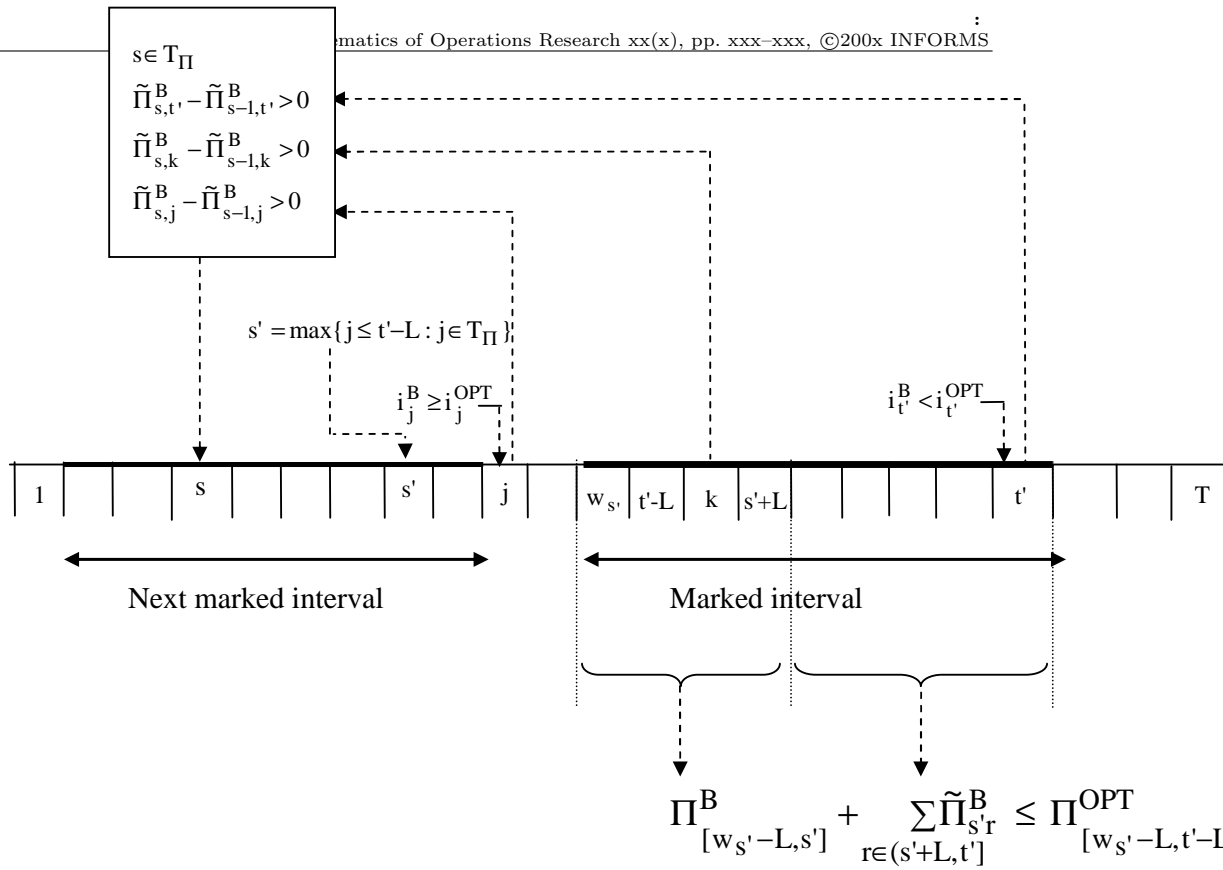


Figure 6: Period s is associated with lost-sales costs from periods j , k and t' . Period j is unmarked, thus, the lost sales costs of OPT can pay for all the contributions of period j to other periods. Periods k and t' are within a marked interval, thus, the lost-sales costs of OPT over the interval can pay for all the contributions to periods in $[1, s']$.

$E[\tilde{\Pi}_t^B | f_t] \leq E[\tilde{\Pi}_t^B | (f_t, W_t = j)]$ for each $t = 1, \dots, T - L$ and $j \in (t, t + L]$. Thus, we get the following theorem.

THEOREM 5.1 *The dual-balancing policy for the capacitated lost-sales model with independent demands, time-dependent holding cost parameters and non-increasing ordering and lost-sales penalty parameters has a worst-case guarantee of 2.*

Finally, we note that the above analysis still holds if the capacities, in each period are random, and follow an exogenous stochastic process. (The only difference is that now the expectations are taken with respect to the random future demands and capacities.)

5.3 Stochastic Non-crossing Lead Times In this subsection, we shall show how the dual-balancing policy and its worst-case analysis can be extended to the lost-sales model with stochastic lead times. Specifically, we assume that the stochastic lead times are such that orders do not *cross* or overtake each other; that is, for every realization of lead times, the order placed in a period arrives no later than any subsequent order. Such lead time models have been first introduced by Kaplan [3] and subsequently streamlined by Nahmias [12]. For ease of exposition we consider the uncapacitated model.

For each $t = 1, \dots, T$, let L_t be the stochastic lead time of the order placed in period t , which is realized at time period $t + L_t$. The stochastic process $\{L_t\}_t$ has the property that $\{t + L_t\}$ is a non-decreasing sequence for every sample path (i.e., with probability 1). We also assume that the processes $\{L_t\}_t$ and $\{D_t\}_t$ are independent of each other.

We will now explain how the definitions, results and proofs discussed in Section 3 above have to be modified to extend the dual-balancing policy and the analysis to the more general model with stochastic lead time.

The dynamics of the inventory-on-hand process is now reflected by the following modified version of

equation (3).

$$I_{t+1}^P = (I_t^P - D_t)^+ + \sum_{r:r+L_r=t+1} Q_r^P. \quad (22)$$

Next, the marginal holding cost incurred by the units ordered in period s , denoted again by H_s^P , is still given by equation (4) after replacing L with L_s , the stochastic lead time of an order placed in s .

Note that, in the model with stochastic lead times, there are no orders delivered in the time interval $(t + L_t, t + 1 + L_{t+1})$ even if $t + L_t < t + 1 + L_{t+1}$. The lower and upper ends of this interval are the periods in which the respective orders placed in periods t and $t + 1$ are delivered. In particular, if $t + L_t < t + 1 + L_{t+1}$, then once the decision of how many units to order in period t is made, the lost-sales costs over $[t + L_t, t + 1 + L_{t+1})$ are just a function of future demands and are not affected by any future decision. Thus, we define $\tilde{\Pi}_s^P$ to be the marginal lost-sales costs associated with period t ,

$$\tilde{\Pi}_t^P = p \sum_{r \in [t+L_t, t+1+L_{t+1})} (D_r - I_r^P)^+. \quad (23)$$

Observe that if the orders placed in t and $t + 1$ arrive in the same period, the marginal lost-sales cost associated with period t is zero. In particular, the lost-sales costs incurred over the planning horizon are associated with periods for which the order placed in the period *strictly precedes* the order in the subsequent period. Let $\mathcal{T}^* = \{t : t + L_t < t + 1 + L_{t+1}\}$. Note that $\tilde{\Pi}_t^P$ is zero if t does not belong to \mathcal{T}^* . The functions, l_t^B and $\tilde{\pi}_t^B$ are as defined above in Section 3 (note that here we deviate from the convention in that l_t^B is not the realization of L_t). As before, in each period s , conditioning on the observed information set f_s , we compute the balancer q_s^B to balance l_s^B against π_s^B . Note that the functions l_s^B and π_s^B are more complex to evaluate in the presence of stochastic lead times. In particular, the corresponding expectations in the definitions of l_s^B and π_s^B are now taken with respect to the future demands *and* the conditional distribution of the lead times of all the outstanding orders. (The information set f_s induces a conditional distributions of the stochastic lead times of all the outstanding orders.) Computing the functions l_s^B and π_s^B highly depends on the complexity of the distributions of the lead times. However, there are many scenarios, in which these functions can still be estimated accurately, for example, using Monte Carlo integration.

Defining the random variable Z_t (for each $t = 1, \dots, T$) in the same way as before, it is readily verified that Lemma 3.1 still holds, i.e., $E[\mathcal{C}(B)] = 2 \sum_t E[Z_t]$.

Let us now revisit the definition of the truncated inventory position, Y_{st}^P . In the model with a deterministic lead time, it includes the inventory on-hand in period s and all orders outstanding at s that were placed by time period t . In particular, $Y_{t+L,t}^P$ is the inventory on-hand at $t + L$ after receiving the order placed in t . In the model with stochastic lead times, it is possible that the order placed in t arrives along with orders placed in subsequent periods. Thus, we redefine $Y_{t+L_t,t}^P$ to be equal to the on-hand inventory at the end of period $t + L_t - 1$, i.e., $(I_{t+L_t-1}^P - D_{t+L_t-1})^+$, plus all the orders delivered in period $t + L_t$ that *were placed by time period t* , i.e., $\sum_{j \leq t: j+L_j=t+L_t} Q_j^P$. That is,

$$\begin{aligned} Y_{st}^P &= I_s^P + \sum_{j \leq t: j+L_j > s} Q_j^P \quad \forall s \in [t, t + L_t) \text{ and} \\ &= (I_{t+L_t-1}^P - D_{t+L_t-1})^+ + \sum_{j \leq t: j+L_j=t+L_t} Q_j^P \text{ for } s = t + L_t. \end{aligned}$$

Observe that under the above definition, the truncated inventory $Y_{t+L_t,t}^P$ does not depend on orders placed after period t even if they are delivered in period $t + L_t$.

Using the modified notion of truncated inventory, we define the sets \mathcal{T}_H , \mathcal{T}_Π and the notion of *alternation* in the same way as before (see Equations (7) and (8) above). Lemma 3.2 still holds, that is $H^{OPT} \geq \sum_{t \in \mathcal{T}_H} H_t^B$. The only modification to the proof of Lemma 3.2 above involves comparing $Y_{t+L_t,t}^B$ and $Y_{t+L_t,t}^{OPT}$ instead of comparing I_{t+L}^B and I_{t+L}^{OPT} .

Similarly, the statement of Lemma 3.3 remains unchanged, whereas the proof requires some modifications. We restate the lemma and give detailed arguments, where necessary, to show how the proof can be extended to the stochastic lead time model.

LEMMA 5.2 *The lost-sales penalty incurred by OPT, denoted by Π^{OPT} , is greater than the lost-sales penalty costs of the dual-balancing policy which are associated with periods $t \in \mathcal{T}_\Pi$, i.e., $\Pi^{OPT} \geq \sum_{t \in \mathcal{T}_\Pi} \tilde{\Pi}_t^B$.*

PROOF. First, observe that

$$\sum_{t \in \mathcal{T}_\Pi} \tilde{\Pi}_t^B = \sum_{t \in \mathcal{T}_\Pi \cap \mathcal{T}^*} \tilde{\Pi}_t^B,$$

since $\tilde{\Pi}_t^B = 0$ if $t \notin \mathcal{T}^*$.

We again describe a random partition of the periods into marked and unmarked periods. Consider a realization of the demands d_1, \dots, d_T and lead times l_1, \dots, l_T . Start at T and look for the latest $s \leq T$ with $i_s^B < i_s^{OPT}$ and some $t \in \mathcal{T}_\Pi \cap \mathcal{T}^*$ such that $s = t + l_t$. If such an s does not exist terminate. Otherwise, let $t' = s - l_{t'}$ and let $w_{t'}$ be again the earliest period r in $[t', t' + l_{t'}]$ for which $y_{r,t'}^B \geq y_{r,t'}^{OPT}$ (where now we refer to the modified notion of the truncated inventory position). Let s'_t be the latest unmarked period within $(t' + l_{t'}, t' + 1 + l_{t'+1})$, where initially all periods are unmarked. We mark the periods $[w_{t'}, s'_t]$ and repeat the above with respect to the interval $[1, w_{t'}]$.

It is easy to see that, for each $t \in \mathcal{T}_\Pi \cap \mathcal{T}^*$, for which $t + L_t$ is not marked, then $I_{t+L_t}^B \geq I_{t+L_t}^{OPT}$ (for otherwise it would have been marked). However, this implies that, for each $s \in [t + L_t, S'_t]$ (where S'_t is the latest unmarked period in the interval $[t + L_t, t + 1 + L_{t+1})$, we have $I_s^B \geq I_s^{OPT}$. The reason is that no order arrives within that interval after period $t + L_t$. In particular, this implies that $\tilde{\Pi}_t^B \leq \Pi_{[t, S'_t]}^{OPT}$, where $\Pi_{[t, S'_t]}^{OPT}$ is the lost-sales costs incurred by OPT over the interval $[t + L_t, S'_t]$.

It is again sufficient to show that for each marked interval $[W_{t'}, S_{t'}]$ the lost sales incurred by OPT over the interval are higher than the respective lost sales of the dual-balancing policy over that interval. The arguments are identical to those used in Lemma 3.3. The proof then follows (see Figure 5.7).

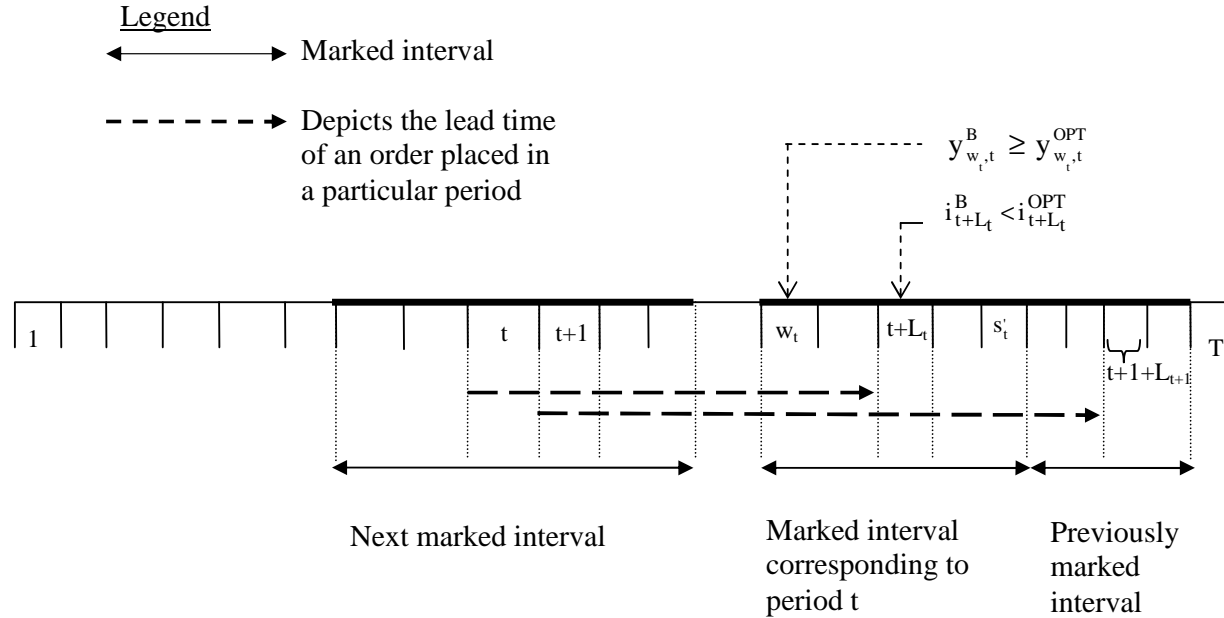


Figure 7: The partition into marked and unmarked periods in the case, where the lead times are stochastic. □

Finally, Equations (14) and (15) still hold. Furthermore, Lemmas 3.4, 3.5 and 3.6 hold (the proofs are almost identical). Thus, we get the following theorem.

THEOREM 5.2 *The dual-balancing policy for the lost-sales model with stochastic (non-crossing) lead times, independent demands, time-dependent holding cost parameters and non-increasing ordering and lost-sales penalty parameters has a worst-case guarantee of 2.*

5.4 Analysis - Correlated Demands In this subsection, we shall show that the analysis in Section 3 above holds in more general settings, where the demands in different periods are correlated. We have already seen that the analysis holds if the Inequalities (16) and (17) hold for each $t = 1, \dots, T - L$ and $f_t \in \mathcal{F}_t$. The following inequalities are equivalent to Inequalities (16) and (17) above. For each period $t = 1, \dots, T - L$ and an information set $f_t \in \mathcal{F}_t$ with $y_{tt}^B < y_{tt}^{OPT}$, we require that

$$E[H_t^B | f_t] \leq E[H_t^B | (f_t, t \in \mathcal{T}_H)] \tag{24}$$

and

$$E[\Pi_t^B | f_t] \geq E[\Pi_t^B | (f_t, t \in \mathcal{T}_H)]. \tag{25}$$

It is clear that Inequality (24) is identical to Inequality (16). In addition, we can write

$$E[\Pi_t^B | f_t] = Pr(t \in \mathcal{T}_H | f_t) \cdot E[\Pi_t^B | (f_t, t \in \mathcal{T}_H)] + Pr(t \in \mathcal{T}_\Pi | f_t) \cdot E[\Pi_t^B | (f_t, t \in \mathcal{T}_\Pi)].$$

However, this and Inequality (25) imply that

$$E[\Pi_t^B | f_t] \leq E[\Pi_t^B | (f_t, t \in \mathcal{T}_\Pi)],$$

which implies Inequality (17).

The main issue in models with correlated demands is that the demands in periods $t + L, \dots, T$ are not independent of the demands over $[t, t + L]$ (we have used this property in the proofs of Lemmas 3.5 and 3.6 for independent demands). Next we consider two natural demand models in which demands in different periods are correlated. We shall show that Inequalities (24) and (25) above hold under these demand structures, thus, the worst-case analysis is valid. Specifically, we consider the *multiplicative auto-regression demand model* in which demands are defined recursively through the formula $D_t = D_{t-1}N_t$ (for each $t = 1, \dots, T$), where $\{N_t\}_t$ are independent lognormal random variables with mean equal to 1. (The assumption that $\{N_t\}_t$ are lognormal is common in the literature, but as noted below is not necessary for the analysis). In addition, we consider the well-known and commonly used *Auto-regression demand model* ($AR(1)$), in which demands are again defined recursively through the formula $D_t = \rho_t D_{t-1} + N_t$ (see, for example, [1]). We assume that $\{\rho_t\}_t$ is a sequence of positive constants and now $\{N_t\}_t$ are independent random variables each with mean equal to 0 (the latter is again without loss of generality).

The analysis for these models is based on a more refined characterization of the event $f_t \cap [t \in \mathcal{T}_H]$ and on the notion of stochastic order between random vectors. Next we characterize the event $f_t \cap [t \in \mathcal{T}_H]$ for cases where f_t , the observed information set at the beginning of period t , is such that $y_{tt}^B < y_{tt}^{OPT}$.

For each $s \in [t, t + L]$, let A_{st} be as before the event $[Y_{st}^B < Y_{st}^{OPT}] \cap [Y_{s+1,t}^B \geq Y_{s+1,t}^{OPT}]$. The event A_{st} implies that in period s there was an *alternation* in the relation between the respective truncated inventory positions of the dual-balancing policy and that of OPT . Similarly, let B_{st} be the event $[Y_{st}^B < Y_{st}^{OPT}] \cap [Y_{s+1,t}^B < Y_{s+1,t}^{OPT}]$, i.e., that in periods s and $s + 1$ the respective truncated inventory of OPT was strictly higher than the one of the dual-balancing policy. It is now clear that we can express

$$f_t \cap [t \in \mathcal{T}_H] = f_t \cap B_{tt} \cap \dots \cap B_{t+L-1,t}.$$

However, in the proof of Lemma 3.4 we have already seen that the event A_{st} can occur only if the cumulative orders placed by the dual-balancing over $[s + 1 - L, t]$ are higher than the cumulative orders placed by OPT over that interval, i.e., $\Delta q_s = \sum_{r=s+1-L}^t q_r^B - \sum_{r=s+1-L}^t q_r^{OPT} \geq 0$ (since we condition on f_t , all the orders placed by the dual-balancing policy and OPT over $(t - L, t]$ are known). We call each of the periods with this property a *potential alternation point*. Let \mathcal{S}_t be the set of all potential alternation points in the interval $[t, t + L]$. The event $f_t \cap [t \in \mathcal{T}_H]$ can then be written as

$$f_t \cap (\cap_{s \in \mathcal{S}_t} B_{st}).$$

Moreover, Lemma 3.4 above implies that this event can be expressed as

$$f_t \cap [t \in \mathcal{T}_H] = f_t \cap (\cap_{s \in \mathcal{S}_t} [D_s \leq I_s^{OPT} - \Delta q_s]). \tag{26}$$

Next we present the analysis for the multiplicative auto-regression demand model. An almost identical analysis can be applied to the $AR(1)$ model. In particular, we use the notion of stochastic order between random vectors. Below we provide a rigorous definition of stochastic order.

DEFINITION 5.1 *Let X and Y be two random n -vectors. We say that X is stochastically larger than Y , denoted as $X \geq_{st} Y$, if $E[\phi(X)] \geq E[\phi(Y)]$ ($E[\phi(X)] \leq E[\phi(Y)]$), for every increasing (decreasing) function $\phi : \mathbb{R}^n \mapsto \mathbb{R}$, for which the above expectations are well-defined.*

Here a decreasing (increasing) function $\phi : \mathbb{R}^n \mapsto \mathbb{R}$ is such that $\phi(x) \geq \phi(y)$ ($\phi(x) \leq \phi(y)$) for each $x \leq y$, where $x, y \in \mathbb{R}^n$ and the inequality $x \leq y$ applies componentwise. The next two theorems provide useful tools to show that two random vectors are stochastically ordered. The theorems are presented without proofs. (To the best of our knowledge, they were first introduced by Veinott [14]. For a comprehensive discussion of stochastic ordering we refer the reader to [11].)

THEOREM 5.3 (Theorem 1.2.1.5 in [11]) *Let X, Y and Z be random variables. If for each value z in the support of Z , we have $(X|Z = z) \leq_{st} (Y|Z = z)$, then $X \leq_{st} Y$.*

THEOREM 5.4 (Theorem 4 in [14] and Theorem 3.3.7 in [11]) *Let X and Y be n -dimensional random vectors. Suppose that $X_1 \leq_{st} Y_1$ and, for each $j = 2, \dots, n$, we have*

$$X_j|(X_1 = x_1, \dots, X_{j-1} = x_{j-1}) \leq_{st} Y_j|(Y_1 = y_1, \dots, Y_{j-1} = y_{j-1}),$$

as long as $x_m \leq y_m$, for each $m = 1, \dots, j-1$. Then $X \leq_{st} Y$.

Recall that in the multiplicative auto-regressive demand model, the demand in period s is defined as $D_s = D_{s-1}N_s$. In particular, conditioning on f_t we know d_{t-1} and, for each $s \geq t$, the demand D_s can be written as $D_s = d_{t-1} \prod_{j=t}^s N_j$. Moreover, conditioning on f_t , the marginal holding cost incurred by the units ordered by the dual-balancing policy in period t can be expressed as a function of d_{t-1} , the observed pipeline vector \bar{p}_t^B and the realization of the random vector $\bar{N}^t = (N_t, \dots, N_T)$. Specifically, let $\phi_t(n_t, \dots, n_T)$ be the marginal holding cost incurred by the units ordered by the dual-balancing policy in period t , conditioning on f_t and on $N_j = n_j$, for each $j = t, \dots, T$. It is clear that ϕ_t is decreasing in the vector $\bar{n}^t = (n_t, \dots, n_T)$. Similarly, let $\psi_t(n_t, \dots, n_T)$ be the lost sales incurred by the dual-balancing policy in period $t+L$, conditioning on f_t and on $N_j = n_j$, for each $j = t, \dots, T$. It is again clear that ψ_t is increasing in $\bar{n}^t = (n_t, \dots, n_T)$. Moreover, $E[H_t^B|f_t] = E[\phi_t(N_t, \dots, N_T)|f_t]$ and $E[H_t^B|(f_t, t \in \mathcal{T}_H)] = E[\phi_t(N_t, \dots, N_T)|(f_t, t \in \mathcal{T}_H)]$. Similarly, $E[\Pi_t^B|f_t] = E[\psi_t(N_t, \dots, N_T)|f_t]$ and $E[\Pi_t^B|(f_t, t \in \mathcal{T}_H)] = E[\psi_t(N_t, \dots, N_T)|(f_t, t \in \mathcal{T}_H)]$. Thus, it is sufficient to show that the vector $\bar{N}^t|f_t = (N_t, \dots, N_T)|f_t$ is stochastically larger than the vector $\bar{N}^t|(f_t, t \in \mathcal{T}_H)$. Definition 5.1 will then imply that $E[\phi_t(N_t, \dots, N_T)|f_t] \leq E[\phi_t(N_t, \dots, N_T)|(f_t, t \in \mathcal{T}_H)]$ and that $E[\psi_t(N_t, \dots, N_T)|f_t] \geq E[\psi_t(N_t, \dots, N_T)|(f_t, t \in \mathcal{T}_H)]$. This implies that the Inequalities (24) and (25) hold and the worst-case analysis follows.

LEMMA 5.3 *Suppose that the demands D_1, \dots, D_T follow the multiplicative auto-regression demand model. Consider some $t = 1, \dots, T-L$ and some information set f_t such that $y_{tt}^B < y_{tt}^{OPT}$. Then the vector $\bar{N}^t|f_t$ is stochastically larger than $\bar{N}^t|(f_t, t \in \mathcal{T}_H)$.*

PROOF. Recall Equation (26) that

$$f_t \cap [t \in \mathcal{T}_H] = f_t \cap (\cap_{s \in \mathcal{S}_t} [D_s \leq I_s^{OPT} - \Delta q_s]),$$

where again \mathcal{S}_t is the set of potential alternation points over the time interval $[t, t+L)$. Assume that $\mathcal{S}_t = \{s_1 < s_2 < \dots < s_m\}$ (the set \mathcal{S}_t is realized by time t). Consider now the sequence of events $A_1 \supseteq A_2 \supseteq \dots \supseteq A_m$, where, for each $j = 1, \dots, m$,

$$A_j = f_t \cap (\cap_{r=1}^j [D_{s_r} \leq I_{s_r}^{OPT} - \Delta q_{s_r}]).$$

Clearly, the event A_m is equivalent to the event $f_t \cap [t \in \mathcal{T}_H]$. Next we show that

$$\bar{N}^t|f_t \geq_{st} \bar{N}^t|A_1 \geq_{st} \bar{N}^t|A_2 \geq_{st} \dots \geq_{st} \bar{N}^t|A_m.$$

First consider the event A_1 . We shall show that $N_t = N_t|f_t \geq_{st} N_t|(f_t, A_1)$. (Observe that N_t is independent of f_t .) Let $Z = (N_{t+1}, \dots, N_T)$, and consider any value in $\bar{n} = (n_{t+1}, \dots, n_T)$ in the support of Z . Given $[Z = \bar{n}]$, the quantity D_{s_1} is increasing in N_t (from zero to infinity) and $I_{s_1}^{OPT}$ is decreasing in N_{t+1} (from a non-negative value to 0). Thus, given $[Z = \bar{n}]$, the event A_1 can be expressed

as $A_1 = f_t \cap [N_t \leq \beta]$ for some $\beta = \beta(n_{t+1}, \dots, n_T)$, where $\beta \in (0, \infty]$. It follows immediately that $N_t|f_t, Z = \bar{n} \geq_{st} N_t|(Z = \bar{n}, A_1)$. Applying Theorem 5.3, we get that $N_t|f_t \geq_{st} N_t|A_1$.

For each $j = t + 1, \dots, s_1$, consider the random variables $N_j|f_t$ and $N_j|A_1$. Next we condition on the events $A = [N_t|f_t = n_t, \dots, N_{j-1}|f_t = n_{j-1}]$ and $A' = [N_t|A_1 = n'_t, \dots, N_{j-1}|A_1 = n'_{j-1}]$ and assume that $n_r \geq n'_r$, for each $r = t, \dots, j - 1$. Using Theorem 5.4 it is sufficient to show that $N_j|(f_t, A) \geq_{st} N_j|(A_1, A')$. However, this can be shown by arguments similar to the one used above, specifically, by conditioning on each value in the support of $Z' = (N_{j+1}, \dots, N_T)|f_t$. We note that $N_j|(f_t, A)$ is identically distributed as $N_j = N_j|f_t$. (Conditioning only on f_t , the random variables N_t, \dots, N_T are assumed to be independent of each other.)

Observe that identical arguments can be applied to show that $\bar{N}^t|A_1 \geq_{st} \bar{N}^t|A_2 \geq_{st} \dots \geq_{st} \bar{N}^t|A_m$. The proof of the lemma then follows. \square

Lemma 5.3 implies the following theorem.

THEOREM 5.5 *Consider the lost-sales model with time-dependent holding cost parameters, non-increasing lost-sales penalty parameters and multiplicative auto regression demand. Then the dual-balancing policy has a worst-case performance guarantee of 2.*

As we have already mentioned, identical arguments can be used to prove an analogous theorem on the AR(1) demand model.

THEOREM 5.6 *Consider the lost-sales model with time-dependent holding cost parameters, non-increasing ordering and lost-sales penalty parameters and AR(1) demand. Then the dual-balancing policy has a worst-case performance guarantee of 2.*

In the more general case, let N_1, \dots, N_T be independent nonnegative random variables. Assume that, for each period s , we have $D_s = g_s(N_1, \dots, N_s)$. Moreover, g_s is an increasing function, and for each $j = 1, \dots, s$, we have $\lim_{N_j \rightarrow \infty} g_s(N_1, \dots, N_s) = \infty$. It can be verified that the proof of Lemma above 5.3 still holds and so is the worst-case analysis.

We believe that there are additional important demand structures for which the worst-performance guarantee can be shown. Providing a general characterization of the properties of the demand structure, required for the analysis to hold, is a very interesting future research direction.

Acknowledgments. This research was conducted while the first author was a postdoctoral fellow at the IBM, T. J. Watson Research Center.

We thank the Associate Editor and the anonymous referees for handful of constructive comments that significantly improved the content and the exposition of this paper.

References

- [1] V. Gaur, A. Giloni, and S. Seshadri. Information sharing in a supply chain under ARMA demand. *Management Science*, 51, 2005.
- [2] S. G. Johansen. Pure and modified base-stock policies for the lost sales inventory system with negligible set-up costs and constant lead times. *International Journal of Production Economics*, 71:391–399, 2001.
- [3] R. Kaplan. A dynamic inventory model with stochastic lead times. *Management Science*, 16(7):491–507, March 1970.
- [4] S. Karlin and H. Scarf. Inventory models of the arrow-harris-marschak type with time lag. In K. J. Arrow, S. Karlin, and H. Scarf, editors, *Studies in the mathematical theory of inventory and production*, chapter 9, pages 155–178. Stanford University Press, 1958.
- [5] R. Levi, M. Pal, R. O. Roundy, and D. B. Shmoys. Approximation algorithms for stochastic inventory control models. Technical Report TR1412, ORIE Department, Cornell University, 2004. To appear in *Mathematics of Operations Research*.
- [6] R. Levi, R. O. Roundy, , and V. A. Truong. Provably near-optimal balancing policies for multi-echelon stochastic inventory control models. Working paper, 2005.
- [7] R. Levi, R. O. Roundy, D. B. Shmoys, and V. A. Truong. Approximation algorithms for capacitated stochastic inventory control models. submitted, 2004.
- [8] T. E. Morton. Bounds on the solution of the lagged optimal inventory equation with no demand backlogging and proportional costs. *SIAM Review*, 11(4):572–596, 1969.

- [9] T. E. Morton. The near-myopic nature of the lagged-proportional-cost inventory problem with lost sales. *Operations Research*, 19(7):1708–1716, 1971.
- [10] A. Muharremoglu and J. N. Tsitsiklis. A single-unit decomposition approach to multi-echelon inventory systems. Working paper, 2001.
- [11] A. Müller and D. Stoyan. *Comparison Methods for Stochastic Models and Risks*. John Wiley & Sons, 2002.
- [12] S. Nahmias. Simple approximations for a variety of dynamic leadtime lost-sales inventory models. *Operations Research*, 27(5):904–924, 1979.
- [13] M. I. Reiman. A new simple policy for a continuous review lost-sales inventory model. Unpublished manuscript, 2004.
- [14] A. F. Veinott. Optimal policy in a dynamic, single product, nonstationary inventory model with several demands classes. *Operations Research*, 13(5):761–778, 1965.
- [15] P. H. Zipkin. *Foundations of inventory management*. The McGraw-Hill Companies, Inc, 2000.
- [16] P. H. Zipkin. Old and new methods for lost-sales inventory systems. Working paper, 2006.
- [17] P. H. Zipkin. On the structure of lost-sales inventory models. Working paper, 2006.