OPERATIONS RESEARCH
Vol. 00, No. 0, Xxxxx 0000, pp. 000–000
ISSN 0030-364X | EISSN 1526-5463 | 00 | 0000 | 0001

INFORMS
DOI 10.1287/xxxx.0000.0000

Approximation Algorithms for Capacitated Stochastic Inventory Control Models

Retsef Levi
Sloan School of Management, MIT, Cambridge, MA, 02139., retsef@mit.edu,

Robin O. Roundy
School of ORIE, Cornell University, Ithaca, NY 14853, robin@orie.cornell.edu,

David B. Shmoys
School of ORIE and Dept. of Computer Science, Cornell University, Ithaca, NY 14853., shmoys@cs.cornell.edu,

Van Anh Truong
School of ORIE, Cornell University, Ithaca, NY 14853, truong@orie.cornell.edu,

We develop the first algorithmic approach to compute provably good ordering policies for a multi-period, capacitated, stochastic inventory system facing stochastic non-stationary and correlated demands that evolve over time. Our approach is computationally efficient and guaranteed to produce a policy with total expected cost no more than twice that of an optimal policy. As part of our computational approach, we propose a novel scheme to account for backlogging costs in a capacitated, multi-period environment. Our cost-accounting scheme, called the *forced marginal backlogging cost-accounting scheme*, is significantly different from the period-by-period accounting approach to backlogging costs used in dynamic programming; it captures the long-term impact of a decision on system performance in the presence of capacity constrains. In the likely event that the per-unit order costs are large compared to the holding and backlogging costs, a transformation of cost parameters yields a significantly improved guarantee. We also introduce new semi-myopic policies based on our new cost-accounting scheme to derive bounds on the optimal base-stock levels. We show that these bounds can be used to effectively improve *any* policy. Finally, empirical evidence is presented that indicates that the typical performance of this approach is significantly stronger than these worst-case guarantees.

Subject classifications: Stochastic Inventory Control; Heuristics; Approximation Algorithms.
Area of review: Supply Chain Management.
History: Submitted August 2005. Revised July 2006, December 2006, May 2007, September 2007, September 2007.

1. Introduction

The periodic-review, capacitated inventory control problem for systems facing stochastic, non-stationary (time-dependent) demands that are correlated and evolve over time is an important classical problem that is widely recognized to be computationally challenging. We develop a new algorithmic approach to compute the order quantity for such a system. We build on the work of Levi et al. (2007), who used a marginal holding cost accounting scheme and cost balancing techniques to derive the first policies with worst-case performance guarantees for uncapacitated models. In this paper, we introduce a novel marginal backlogging cost accounting scheme that, in combination with their techniques, lead to analogous results for the much harder capacitated model. We believe that our new cost accounting scheme will have applications in many other settings. Our algorithm is guaranteed to compute a solution of total expected cost no more than twice that of an optimal policy for any instance of the problem. The algorithm is computationally efficient and implementable without having to enumerate exhaustively future scenarios and corresponding future decisions. In particular, the decision made in the current period is unaffected by any future decision. Thus, it can be implemented efficiently even in the presence of complex demand structures.

Specifically, we consider single-item models with one location and a finite planning horizon of T discrete time periods. The demands over the T periods are random variables that can be non-stationary and correlated. The costs consist of a per-unit, time-dependent ordering cost, a per-unit holding cost for carrying excess inventory from period to period and a per-unit backlogging cost, which is a penalty incurred, in each period, for each unit of unsatisfied demand (where all shortages are fully backlogged). There is a time-dependent capacity constraint on the number of units ordered in each period and a lead time between the time that an order is placed and the time that it actually arrives. The capacity constrains and lead times may be stochastic.

Capacitated problems are inherently more difficult computationally compared to their uncapacitated counterparts. The constraint on capacity makes future costs heavily dependent on current decisions. Myopic policies, which do not consider the impact of a decision made in the current period on the costs incurred in future periods, seem to perform well for some scenarios in uncapacitated systems and are even optimal in some settings (seeVeinott (1965), Ignall and Veinott (1969) Iida and Zipkin (2001), and Lu et al. (2006)). However, when applied to capacitated problems, they usually perform very poorly because they do not consider possible capacity limitations in future periods.

In this work, we introduce a look-ahead backlogging cost-accounting scheme, called the *forced marginal backlogging cost-accounting scheme*, to capture the long-term impact of current decisions on future costs in the presence of capacity constrains. Our new cost accounting scheme assigns to the decision in each period *all* of the expected backorder costs that, once this decision is made, become inevitable; that is, they are unaffected by any decision made in future periods, and are dependent only on future demands. The forced marginal backlogging cost reduces to the traditional backlogging cost when the capacity is infinite; thus, it is a generalization of the traditional backlogging cost. Finally, as discussed in Section 3.1, it is straightforward to compute in most common scenarios.

The key feature distinguishing the algorithms presented in this paper from those previously studied for capacitated systems is the treatment of correlation in demand across time as well as non-stationarity. Moreover, we allow observations of the past to change demand forecasts for the future. Our model also captures other important characteristics of a non-stationary environment: the parameters are fully time-dependent, including cost parameters and system capacity. An important application of demand correlation and non-stationarity is in the use of dynamic demand forecasts. These forecasts and the way they evolve over time provide vital information that can be used to find effective inventory control policies in dynamic and highly volatile demand environments. The assumptions that we make on the demand distributions in this work are mild enough to generalize all of the currently known approaches in the literature to model correlation and non-stationarity of demand over time. These include classical approaches like the *martingale model of forecast evolution model* (MMFE), exogenous Markovian demand, time series, order-one auto-regressive demand and random walks. For an overview of the different approaches and models, and for relevant references, we refer the reader to Iida and Zipkin (2001) and Dong and Lee (2003).

High correlation between demands across different periods in non-stationary and dynamic environments presents a considerable challenge to computing, or even approximating, optimal inventory control policies. The dominant paradigm in almost all of the existing literature has been to formulate multi-period capacitated models using dynamic programming. The optimization problem is defined recursively over time using subproblems for each possible state of the system. The state usually consists of a given time period, the level of inventory at the beginning of the period, the resulting conditional distribution of future demands over the rest of the horizon, and possibly more information that is available by that period. For each subproblem, an optimal solution is computed to minimize the expected overall discounted cost from the current point to the end of the horizon. This framework has turned out to be very effective in characterizing the structure of the optimal policy of the overall system. Assuming stationary linear costs and independent and identically distributed (i.i.d.) demands, Federgruen and Zipkin (1986a,b) showed that a modified, base-stock policy is optimal under infinite-horizon average-cost and discounted cost criteria. They established the existence of a target inventory level such that the optimal policy aims to keep inventory levels as close as possible to that target. When the inventory level at the beginning of the period is above the target level, the optimal policy does not order. When the inventory level at the beginning of the period is below the target level, it might not be possible to order up to the target level because of the capacity constraint. In this case, the order placed would be up to capacity. Tayur Kapuscinski and Tayur (1998) and Aviv and Federgruen (1997) derived the optimal policy in the same settings, but for independent cyclical demands.

Axsäter (1990) is the first to introduce the notion of *matching* between pairs of demand and supply units. Specifically, he observes that in a distribution system with a single depot and multiple retailers, a supply unit ordered by a retailer can be used to fill a corresponding demand unit following a certain order. He matches this pair of units and evaluates the corresponding expected holding cost. Katircioglu and Atkins (1996) have used this observation to analyze the optimal policies in unit demand inventory systems. For the uncapacitated periodic-review stochastic inventory control problem, Muharremoglu and Tsitsiklis (2001) have proposed an alternative approach to the dynamic programming framework. They have observed that this problem can be decoupled into a series of *unit supply-demand subproblems*, where each subproblem corresponds to a single unit of supply and a single unit of demand that are matched. This novel approach enabled them to substantially simplify some of the dynamic programming based proofs on the structure of optimal policies, as well as to prove several important new structural results. In particular, they have established the optimality of state-dependent base-stock policies for the uncapacitated model with general Markov-modulated demand. Using this unit decomposition, they have also suggested new methods to compute the optimal policies. However, their computational methods are essentially dynamic programming approaches applied to the unit subproblems, hence they suffer from similar problems in the presence of correlated and non-stationary demand. Although our approach is very different from theirs, we use some of their ideas as technical tools in some of the proofs. Janakiraman and Muckstadt (2003) have extended this approach to capacitated models and established the optimality of state-dependent modified base-stock policies for models with Markov-modulated demand.

Unfortunately, the rather simple forms of these policies do not always lead to efficient algorithms for computing the optimal policies. Complex demand structures, such as the one we consider in this work, cause the state space of the corresponding dynamic programs to explode (see Iida and Zipkin (2001), and Dong and Lee (2003) for relevant discussions on the MMFE model). There does not exist at present, nor is there likely to be developed, an efficient algorithm to solve these dynamic programs to optimality, even for the uncapacitated model. The difficulty comes from the fact that we need to solve 'too many' subproblems, a phenomenon known as *the curse of dimensionality*. To date, computational procedures have been made tractable only under assumptions of simple demand structures. If the demands in different periods are independent, the corresponding dynamic programs are relatively straightforward to solve. Dynamic programming can still be tractable for uncapacitated models with Markov-modulated demand but under rather strong assumptions on the structure and the size of the state space of the underlying Markov process (see, for example, Song and Zipkin (1993) and Chen and Song (2001)). Tayur (1992) uses the shortfall distribution and the theory of storage processes to derive an efficient computational method for computing the optimal policy in the stationary cost, i.i.d. demand, average-cost case. Roundy and Muckstadt (2000) showed how to obtain approximate base stock levels, also for the stationary cost and i.i.d. demand case, by approximating the distribution of the shortfall process. Kapuscinski and Tayur (1998) proposed a simulation-based technique using infinitesimal perturbation analysis to compute the optimal policy for capacitated problems with independent, cyclical demands.

There have been heuristic approaches to compute order quantities for capacitated problems. However, we are aware of very few attempts to analyze the worst-case performance of heuristics and most bounds derived are dependent on the particular input (see, for example, Lu et al. (2006)). To the best of our knowledge, there are no other policies for stochastic inventory control models with constant worst-case performance guarantees. Metters (1997) found heuristics for capacitated, lost-sales problems with independent, cyclical demands. Chan (1999) have considered heuristics for uncapacitated and capacitated multi-item models. Instead of solving the one-period problem (as in the Myopic policy), they have added a penalty function to the one-period problem, which they called the Q-function. This function accounts for the holding cost incurred by the inventory left at the end of the period over the entire horizon. Their look-ahead approach with respect to the holding cost is somewhat related to our approach, though significantly different.

As we have already mentioned, this paper builds on the work of Levi et al. (2007). They give the first algorithms with a constant performance guarantee for the uncapacitated stochastic inventory control model with correlated, non-stationary demands; specifically, their algorithms always find solutions of total expected cost no more than twice the optimal. Their algorithms are based on two main ideas. First, they construct a look-ahead holding cost accounting scheme, called the *marginal holding cost accounting scheme*, to compute the additional holding costs incurred by units ordered in the current period throughout the entire horizon. Secondly, they use *cost-balancing* techniques in that, in each period, they order exactly to balance the following two opposing costs: the conditional expected marginal holding cost against the conditional expected period backlogging cost a lead time ahead. Their approach relies heavily on the ability of the system to order in each period a 'balancing quantity' that equalizes the expected marginal holding cost and the expected backlogging cost in the period. In capacitated systems, the approach fails because this balancing quantity might not be attainable due to capacity constrains. Our forced marginal backlogging cost accounting scheme is designed to remedy this problem by reassigning backlogging costs more appropriately to the decisions that create them, enabling us to find a 'balancing order quantity' for capacitated systems. Suppose that in the current period the order placed was not up to capacity; we wish to account for the potential backlogging cost in future periods incurred directly by the decision not to use the full available capacity. Assume temporarily that we order up to capacity in each one of the periods. Suppose now that in the current period we do not order up to capacity. Then the expected marginal backlogging cost associated with the current period is the overall increase in the expected backlogging cost over the entire horizon resulting from this decision. In this way, our balancing policy for a capacitated system is able to achieve the same worst-case performance guarantee of 2, with surprisingly little additional computational effort. When applied to uncapacitated models the policies described in this paper are identical to the Dual-Balancing policies described by Levi et al. (2007). Thus, they can be viewed as generalizations of the original Dual-Balancing policies to capacitated inventory models.

We also use the marginal holding and forced marginal backlogging cost accounting schemes to derive additional semi-myopic policies, called the *Lower-Myopic* and *Upper-Myopic* policies. The policies provide lower and upper bounds on the optimal base-stock levels, respectively, which can be used in conjunction with *any* policy to achieve lower expected cost.

Furthermore, in Section 4.2 we show how to use standard cost transformations to improve the performance of the algorithms in many important settings (see also Levi et al. (2007)). These transformations yield a modified instance of the problem that is equivalent to the original one from an optimization perspective, but models only holding and backlogging costs. If the per-unit ordering cost is constant over time, then applying our algorithms to the modified instance yields an approximation algorithm with a worst-case guarantee of 2 with respect to the holding and backlogging costs, and which has the same total per-unit ordering cost as the optimal policy. More generally, when the ordering costs are large, the worst-case performance guarantee of the modified-cost Dual-Balancing policy will be much better than 2.

In Section 6 we test the typical performance of the balancing algorithms in two settings. We consider an inventory system that has i.i.d. demand (no correlations), and a demand distribution with an exponential tail, because the optimal policy can be computed analytically. (The motivation is to the test balancing policies at least in one environment, in which the optimal policy and cost are known.) However, balancing policies are most attractive in scenarios with complex demand structures, whereas optimal policies can not be computed and no provable good heuristics or reasonable lower bounds are known. Thus,we also consider the same set of test scenarios tested in Hurley et al. (2006), in which the uncapacitated versions of these algorithms were evaluated computationally. In these scenarios the demands and forecasts evolve according to the multiplicative MMFE model. Optimal policies are not computable and strong lower bounds do not exist, so we used the Myopic policy as a benchmark for evaluating performance. The performance of the Balancing policies is very robust. It was within 11% of optimal on average in the first test (always within 25%), and consistently improved upon myopic, by over 27%, on average and by over 50% in many scenarios.

The paper is organized as follows. In Section 2 we present the mathematical formulation of the periodic-review, capacitated, stochastic inventory control problem. In Section 3 we describe the forced marginal backlogging cost accounting scheme for the capacitated model. In Section 4 we describe the balancing policy and its worst-case analysis. We also extend the approximation results to the case of discrete demand and stochastic lead times (see Appendix C). In Section 5, we develop lower and upper bounds on the optimal inventory levels, and show how to use them to improve any policy. Section 6 contains our computational results. Appendix A contains a very simple, illustrative example for the case of integer-valued demand. In Appendix B we present a detailed description of the scenarios tested in the computational results.

2. Capacitated Periodic-Review Stochastic Inventory Control Problem

In this section, we provide the mathematical formulation of the capacitated periodic-review stochastic inventory problem 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 unlike the convention above). The demands over these periods are random variables, denoted by $D_1,\dots,D_T$.

As part of the model, we assume that at the beginning of each period s, we are given what we call an *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 $[1,s)$, and possibly some more (external) information denoted by $(w_1,\dots,w_s)$. The information set f_s in period s is one specific realization in the set of all possible realizations of the random vector $F_s=(D_1,\dots,D_{s-1},W_1,\dots,W_s)$. This set is denoted by $\mathcal{F}_s$. In addition, we assume that in each period s there is a known conditional joint distribution of the future demands $(D_s,\dots,D_T)$, denoted by $I_s:=I_s(f_s)$, which is determined by f_s (i.e., knowing f_s, we also know $I_s(f_s)$). For ease of notation, D_t will always denote the random demand in period t according to the conditional joint distribution I_s for some $s\leq t$, where it will be clear from the context to which period s we refer. We will use t as the general index for time, and s will always refer to the current period. The only assumption on the demands is that for each $s=1,\dots,T$, and each $f_s\in\mathcal{F}_s$, the conditional expectation $E[D_t|f_s]$ is well defined and finite for each period $t\geq s$. In particular, we allow non-stationarity and correlation between the demands of different periods.

In the periodic-review stochastic inventory control problem, our goal is to supply each unit of demand while attempting to avoid ordering it either too early or too late. In period t, $(t=1,\dots,T)$ three types of costs are incurred, a per-unit ordering cost c_t for ordering up to u_t units, where

$u_t \geq 0$ is the available order capacity in period t, a per-unit holding cost h_t for holding excess inventory from period t to $t+1$, and a per-unit backlogging penalty p_t that is incurred for each unsatisfied unit of demand at the end of period t. Unsatisfied units of demand are usually called *backorders*. Backorders fully accumulate over time until they are satisfied. That is, each unit of unsatisfied demand will stay in the system and will incur a backlogging penalty in each period until it is satisfied. In addition, there is a lead time of L periods between the time an order is placed and the time at which it actually arrives. We first assume that the lead time is a known integer L. In Appendix C, we show that our policy can be modified to handle stochastic lead times under the assumption of no order crossing (i.e., any order arrives no later than those placed later in time). In Section 4.1, we show that extensions to the case of random capacities are straightforward.

There is also a discount factor $\alpha \leq 1$. The cost incurred in period t is discounted by a factor of α^t. Since the horizon is finite and the cost parameters are time-dependent, we can assume without loss of generality that $\alpha = 1$. We also assume that there is no speculative motivation for holding inventory or having back orders in the system. To enforce this, we assume that, for each $t = 2, \ldots, T-L$, the inequalities $c_t \leq c_{t-1} + h_{t+L-1}$ and $c_t \leq c_{t+1} + p_{t+L}$ are maintained (where $c_{T+1} = 0$). (If there is a discount factor, we require that $\alpha c_t \leq c_{t-1} + \alpha^L h_{t+L-1}$ and $c_t \leq \alpha c_{t+1} + \alpha^L p_{t+L}$). We also assume that the parameters h_t, p_t and c_t are all non-negative. Note that the parameters h_T and p_T can be defined to take care of excess inventory and back orders at the end of the planning horizon. In particular, p_T can be set sufficiently high so as to ensure that there are very few back orders at the end of period T.

The goal is to find a feasible ordering policy (i.e., one that respects the capacity constraints) that minimizes the overall expected discounted ordering cost, holding cost and backlogging cost. 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 and the current inventory level.

Throughout the paper we will use $D_{[s,t]}$ to denote the accumulated demand over the interval $[s,t]$, i.e., $D_{[s,t]} := \sum_{j=s}^{t} D_j$. We will also use superscripts P and OPT to refer to a given policy P and the optimal policy respectively.

Given a feasible policy P, we describe the dynamics of the system using the following terminology. We let NI_t denote the *net inventory* at the end of period t, which can be either positive (in the presence of physical on-hand inventory) or negative (in the presence of back orders). Since we consider a lead time of L periods, we also consider the orders that are on the way. The sum of the units included in these orders, added to the current net inventory is referred to as the *inventory position* of the system. We let X_t be the inventory position at the beginning of period t *before* the order in period t is placed, i.e., $X_t := NI_{t-1} + \sum_{j=t-L}^{t-1} Q_j$ (for $t = 1, \ldots, T$), where Q_j denotes the number of units ordered in period j (we will sometime denote $\sum_{j=t-L}^{t-1} Q_j$ by $Q_{[t-L,t-1]}$). Similarly, we let Y_t be the inventory position *after* the order in period t is placed, i.e., $Y_t = X_t + Q_t$. Note that once we know the policy P and the information set $f_s \in \mathcal{F}_s$, we can easily compute ni_{s-1}, x_s and y_s, where again these are the realizations of NI_{s-1}, X_s and Y_s, respectively.

3. Marginal Cost Accounting Scheme

In this section, we present a *marginal cost accounting* for stochastic inventory control problems with capacity constraints on the size of the order in each period. This extends and generalizes the marginal cost accounting discussed by Levi et al. (2007). Since this cost accounting approach is central for our approximation results, we explain it in detail, repeating some of the ideas of that paper. Our approach differs significantly from the traditional cost accounting approaches, which is based on standard dynamic programming.

We start by reviewing their cost accounting approach, which is called *marginal cost accounting*. The main idea underlying this approach is to account for *all* the expected costs associated with the decision of how many units to order in period t when this decision is made. More specifically, the decision in period t is associated with all the expected cost that, after that decision is made, become unaffected by any future decision, and are only dependent on future demands. In Levi et al. (2007) it was shown that in uncapacitated models, these costs are relatively easy to compute already in period t, even though they may include costs that are going to be incurred only in future periods. Taking this approach, Levi, Pál, Roundy and Shmoys have proposed a marginal holding cost accounting scheme. Their approach is based on the convention that units in inventory are consumed on a first-ordered-first-consumed basis. This implies that the overall holding cost of the q_s units ordered in period s (i.e., the holding cost they incur over the entire horizon $[s,T]$) is a function only of future demands, and is independent of any future decision. Based on the assumption that inventory is consumed on a first-ordered-first-consumed basis, the q_s units on order will be used to satisfy demand only when the x_s units presently in the system have been completely consumed. Among these q_s units, the number of those still remaining in inventory at the end of period j (where $j \geq s+L$) is precisely $(q_s - (D_{[s,j]} - x_s)^+)^+$. Each of these units incurs a cost of h_j. More specifically, conditioning on an information set $f_s \in \mathcal{F}_s$, the marginal holding cost is defined to be (assuming again that $\alpha = 1$) $\sum_{j=s+L}^{T} h_j(q_s - (D_{[s,j]} - x_s)^+)^+$. Observe again that for each non-anticipatory policy P, if conditioned on some $f_t \in \mathcal{F}_t$, the inventory position at the beginning of period t, denoted by x_t^P, is known deterministically. In addition, once the order in period s is determined, the backlogging cost a lead time ahead in period $s+L$, i.e., $p_{s+L}(D_{[s,s+L]} - (x_s + q_s))^+$, is also dependent only on the future demands. This leads to a marginal cost accounting. For each feasible policy P, let H_t^P be the ordering and holding cost incurred over the interval $[t,T]$ by the Q_t^P units ordered in period t (for $t = 1,\dots,T$), and let Π_t^P be the backlogging cost incurred a lead time ahead in period $t+L$ ($t = 1-L,\dots,T-L$). That is, $H_t^P = c_t Q_t^P + \sum_{j=t+L}^{T} h_j(Q_t^P - (D_{[t,j]} - X_t^P)^+)^+$ and $\Pi_t^P := p_{t+L}(D_{[t,t+L]} - (X_{t+L}^P + Q_t^P))^+$ (where $D_j := d_j$ with probability 1 and $Q_j^P = q_j$ is given as an input for each $j \leq 0$). Let $\mathcal{C}(P)$ be the cost of the policy P. Clearly,

$$\mathcal{C}(P) := \sum_{t=1-L}^{0} \Pi_t^P + H_{(-\infty,T]} + \sum_{t=1}^{T-L} (H_t^P + \Pi_t^P), \tag{1}$$

where $H_{(-\infty,T]}$ denotes the total expected holding cost incurred over the interval $[1,T]$ by units ordered before period 1. We note that the first two expressions $\sum_{t=1-L}^{0} \Pi_t^P$ and $H_{(-\infty,T]}$ are not affected by our decisions (i.e., they are the same for any feasible policy and each realization of the demands). Note that, without loss of generality, we can assume that $Q_t^P = H_t^P = 0$ for any policy P and each period $t = T-L+1,\dots,T$, since nothing that is ordered in these periods can be used within the given planning horizon.

In models with no capacity constraints there is a fundamental difference between holding cost and backlogging cost. In particular, any mistake of ordering 'too little' can be fixed in the next period to avoid further backlogging cost. In particular, the decision of how many units to order affects the backlogging cost in a single period. However, the effect of this decision, if we have ordered 'too much', may last for a number of periods depending on the realized future demands. That is, no future decision can fix this mistake, since we can not order a negative quantity. Consequently, Π_t^P only accounts for costs incurred in a single period, namely, backlogging cost in period $t+L$, and H_t^P accounts for holding costs incurred over multiple periods. By way of contrast, in models with capacity constraints on the size of the order in each period, the above observation is no longer valid. More specifically, because of the capacity constraints, it is no longer true that a mistake of ordering 'too little' in the current period can always be fixed by decisions made in future periods.

3.1. Marginal Backlogging Cost Accounting

We now present a new backlogging cost accounting that associates with the decision of how many units to order in period s what we shall call *forced backlogging cost* resulting from this decision in future periods.

Consider some period s. Suppose that x_s is the inventory position at the beginning of period s and that the number of units ordered in the period is $q_s < u_s$. Let $\bar{q}_s$ be the resulting *unused slack capacity* in period s, i.e., $\bar{q}_s = u_s - q_s > 0$. Focus now on some future period $t \geq s+L$ when this order arrives and becomes available. Suppose that for some realization of the demands, we have that $d_{[s,t]} - (x_s + q_s + \sum_{j\in(s,t-L]} u_j) > 0$. This implies that there exists a shortage in period t, and moreover, even if in every period *after* period s and until period $t-L$ the orders had been up to the maximum available capacity, this part of the shortage in period t would still exist and incur the corresponding backlogging cost. The actual shortage may be even bigger and equal to $d_{[s,t]} - (x_s + q_s + \sum_{j\in(s,t-L]} q_j) > 0$ (recall that $q_j \leq u_j$ for each period j). In other words, given our decision in period s, this part of the shortage could not be avoided by any decision made over the interval $(s, t-L]$ (clearly, any order placed after period $t-L$ will not be available by time t). We conclude that, if more units had been ordered in period s, then at least some of the shortage in period t could have been avoided. More precisely, the maximum number of units of shortage that could have been avoided by ordering more units in period s is equal to $\min\{\bar{q}_s,\ (d_{[s,t]} - (x_s + q_s + \sum_{j\in(s,t-L]} u_j))^+\}$. The intuition is that by ordering more units in period s, we could have averted part of the shortage in period t, but clearly not more than the unused slack capacity $\bar{q}_s$, since we could not have ordered in period s more than additional $\bar{q}_s$ units. In this case, we would say that this part of the backlogging cost in period t was *forced* by the decision in period s, and hence period s is associated with a backlogging penalty of $p_t \min\{\bar{q}_s,\ (d_{[s,t]} - (x_s + q_s + \sum_{j\in(s,t-L]} u_j))^+\}$. This is significantly different from the *traditional* backlogging cost accounting, in which this cost would be associated with period $t-L$.

We let W_{st} be the shortage in period t that is forced by the decision in period s (where again $s \leq t-L$), i.e.,

$$W_{st} := \min\{\bar{Q}_s,\ (D_{[s,t]} - (X_s + Q_s + \sum_{j\in(s,t-L]} u_j))^+\}.$$

An alternative way to express W_{st}, using $\min(a,(b)^+) = (b)^+ - (b-a)^+$ for $a \in \mathbf{R}_+$ and $b \in \mathbf{R}$, is

$$W_{st} := (D_{[s,t]} - (X_s + Q_s + \sum_{j\in(s,t-L]} u_j))^+ - (D_{[s,t]} - (X_s + \sum_{j\in[s,t-L]} u_j))^+. \tag{2}$$

Now using the equalities, $NI_t = X_s + Q_s + \sum_{j\in(s,t-L]} Q_j - D_{[s,t)}$ (for each $s \leq t-L$) and $u_j = Q_j + \bar{Q}_j$ (for each $j = s,\dots,t-L$), we conclude that equation (2) can be written as

$$(D_t - NI_t - \sum_{j\in(s,t-L]} \bar{Q}_j)^+ - (D_t - NI_t - \sum_{j\in[s,t-L]} \bar{Q}_j)^+. \tag{3}$$

To see why (2) (and hence, (3)) holds, observe that $(D_{[s,t]} - (X_s + Q_s + \sum_{j\in(s,t-L]} u_j))^+ > \bar{Q}_s$ if and only if $(D_{[s,t]} - (X_s + \sum_{j\in[s,t-L]} u_j))^+ > 0$. Next we describe several properties of the parameters W_{st}. Clearly, if $\bar{Q}_s = 0$ (i.e. $Q_s = u_s$), then $W_{st} = 0$ for each $t \geq s+L$. It is also readily verified from (3) that if $W_{st} > 0$ for some $s \leq t-L$, then we have $W_{jt} = \bar{Q}_j$ for each $j \in (s, t-L]$.

For each $s = 1-L,\dots,T-L$, let $\tilde{\Pi}_s^P$ be the overall forced backlogging cost in periods $s+L,\dots,T$ associated with period s, i.e., $\tilde{\Pi}_s^P = \sum_{t=s+L}^{T} p_t W_{st}^P$ (we again assume that $D_j = d_j$ with probability 1 for each $j \leq 0$). Let $u_{-L} = \infty$, $q_{-L} = 0$ and $\bar{q}_{-L} = \infty$, and also define, for each $t = 1,\dots,T$,

$$W_{-L,t} := (D_{[1-L,t]} - (x_{1-L} + \sum_{j\in[1-L,t-L]} u_j))^+ = (D_t - NI_t - \sum_{j\in[1-L,t-L]} \bar{Q}_j)^+,$$

and $\tilde{\Pi}^P_{-L} = \tilde{\Pi}_{-L} := \sum_{t=1}^{T} p_t W_{-L,t}$. The last definition of $\tilde{\Pi}_{-L}$ is meant to account for the forced backlogging cost which is independent of *any* decision, and is forced by the demands on *any* feasible policy. It is now readily verified that, for each $t = 1, \ldots, T$ and for each policy P, we have $\Pi^P_{t-L} = p_t(D_t - NI^P_t)^+ = p_t \sum_{j=-L}^{t-L} W^P_{jt}$ (the sum $\sum_{j=-L}^{t-L} W_{jt}$ is telescopic). This implies the following theorem.

THEOREM 1. *Let P be a non-anticipatory policy. Then the cost of policy P can be expressed as $\mathcal{C}(P) := \sum_{t=-L}^{0} \tilde{\Pi}^P_t + H_{(-\infty,T]} + \sum_{t=1}^{T-L}(H^P_t + \tilde{\Pi}^P_t)$.*

Note that the first two terms of $\mathcal{C}(P)$ in Theorem 1, $\sum_{t=-L}^{0} \tilde{\Pi}^P_t$ and $H_{(-\infty,T]}$, are independent of any decision we make and are common to all feasible policies. Recall that $\sum_{t=-L}^{0} \tilde{\Pi}^P_t$ represents the forced backlogging penalty that is forced on any feasible policy. Since these two terms are also non-negative, we omit them from the analysis. This does not impact our approximation results. From now on, we will write the cost of a feasible policy P as $\mathcal{C}(P) = \sum_{t=1}^{T-L}(H^P_t + \tilde{\Pi}^P_t)$. In Appendix A we provide an illustrative example of the our new cost accounting approach.

The intuition is that once a shortage is incurred in period t, it is allocated to past periods $s \le t-L$ in which the orders were below the available capacity. More specifically, the shortage and the resulting backlogging cost in period t are charged to periods $s \le t-L$ with positive unused slack capacity going backward in time from period $t-L$. Each period $s \le t-L$, can be charged with a part of the backlogging cost in period t for up to $\bar{q}_s$ units, the unused slack capacity in

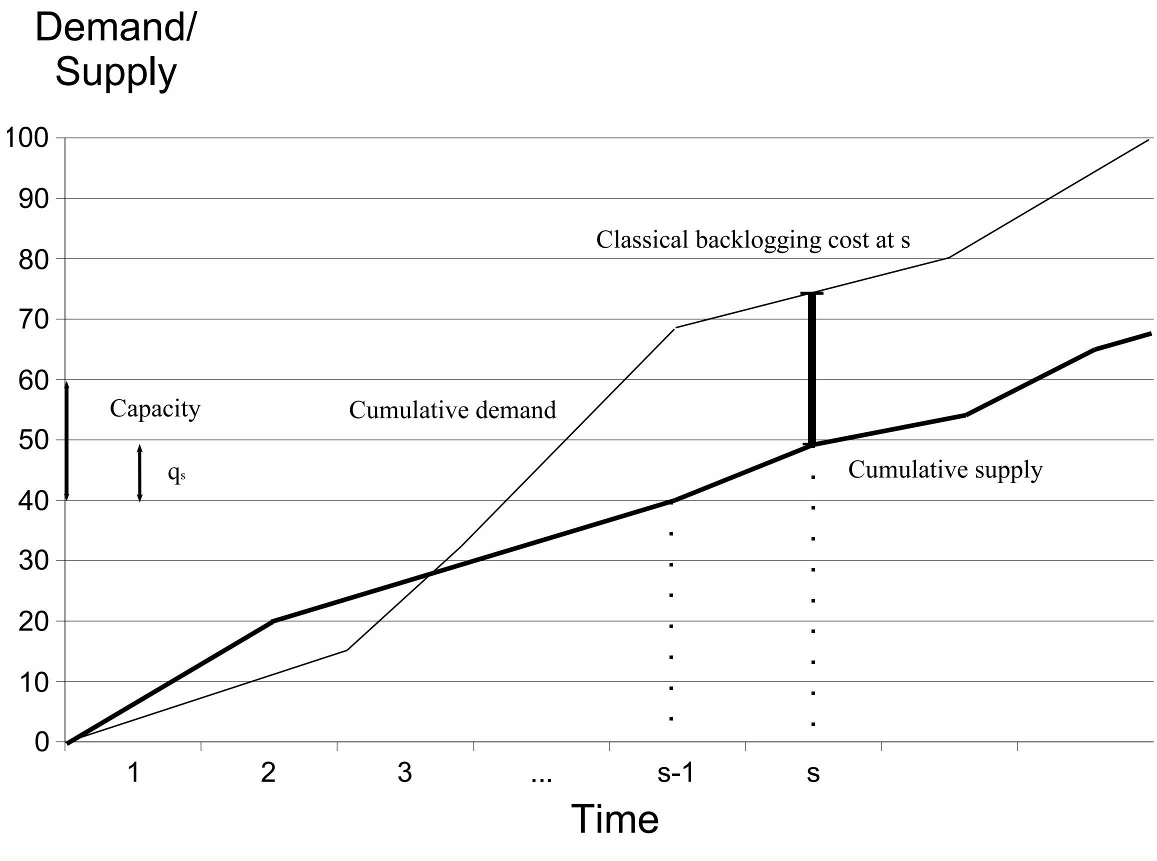

Figure 1 Period-by-period backorder cost accounting

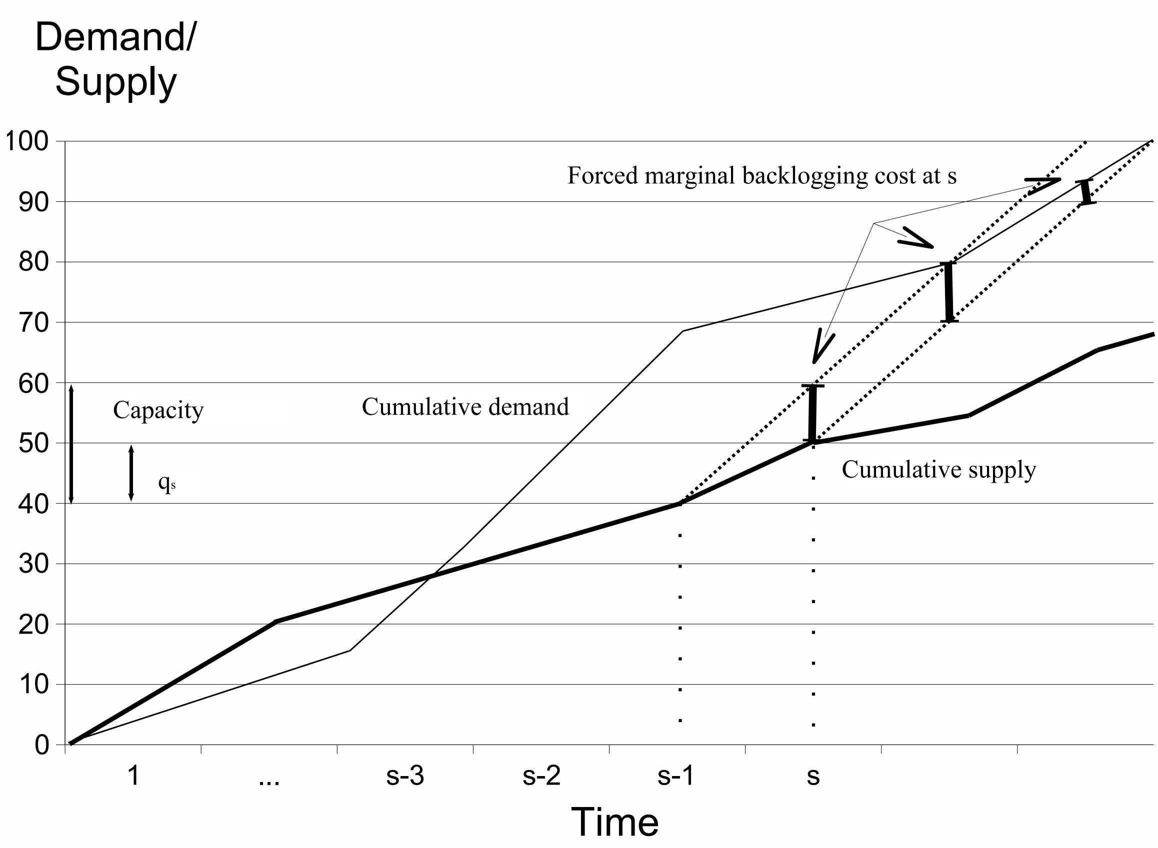

Figure 2 Forced marginal backorder cost accounting

Figures 1, 2 and 3 illustrate graphically the difference between classical period-by-period accounting and forced marginal accounting for backlogging costs. All three figures reflect a single sample path of demands and orders. The total backlogging cost over the horizon is the area above the cumulative supply curve (thick line) and below the cumulative demand curve (thin line). Classical period-by-period accounting assigns to period s the difference between the curves at s (see Figure 1). Forced marginal accounting of backlogging costs assigns to period s all of the backlogging costs that were "forced", or made inevitable, because we did not order to capacity in period s. This corresponds to the area inside of the trapezoid shown in Figure 2. This trapezoid is created by extending the cumulative supply curve, starting at $s-1$ and at s, to the right at a slope equal to the capacity of the system. These lines represent what the supply curves would look like if our policy consistently ordered at full capacity from $s-1$ and s onwards, respectively. In fact, consider the thick short bars in the trapezoid in Figure 2. The first and second terms of (2) are the vertical coordinates of the end points of these bars. Consequently each W_{st}, for $t > s$, is the length of one of these bars. Figure 3 takes a different point of view. It considers the backlogging costs incurred in period s, and illustrates how those costs are allocated to periods $s, s-1, \ldots, 1, \ldots, -L$.

In Levi et al. (2007) it is shown that the marginal holding cost consists of a sum of partial expectations. Once x_s is known at time s, the summands are expectations of simple piecewise linear functions. If the accumulated demand $D_{[s,j]}$ (for each j, s) has any of the distributions that are commonly used in inventory theory (e.g., Normal, Gamma, Lognormal, Laplace, etc) (Zipkin 2000), then it is extremely easy to evaluate these terms. If the distribution of $D_{[s,j]}$ is discrete, these functions can be computed recursively in efficient ways using the CDF functions. More generally, the

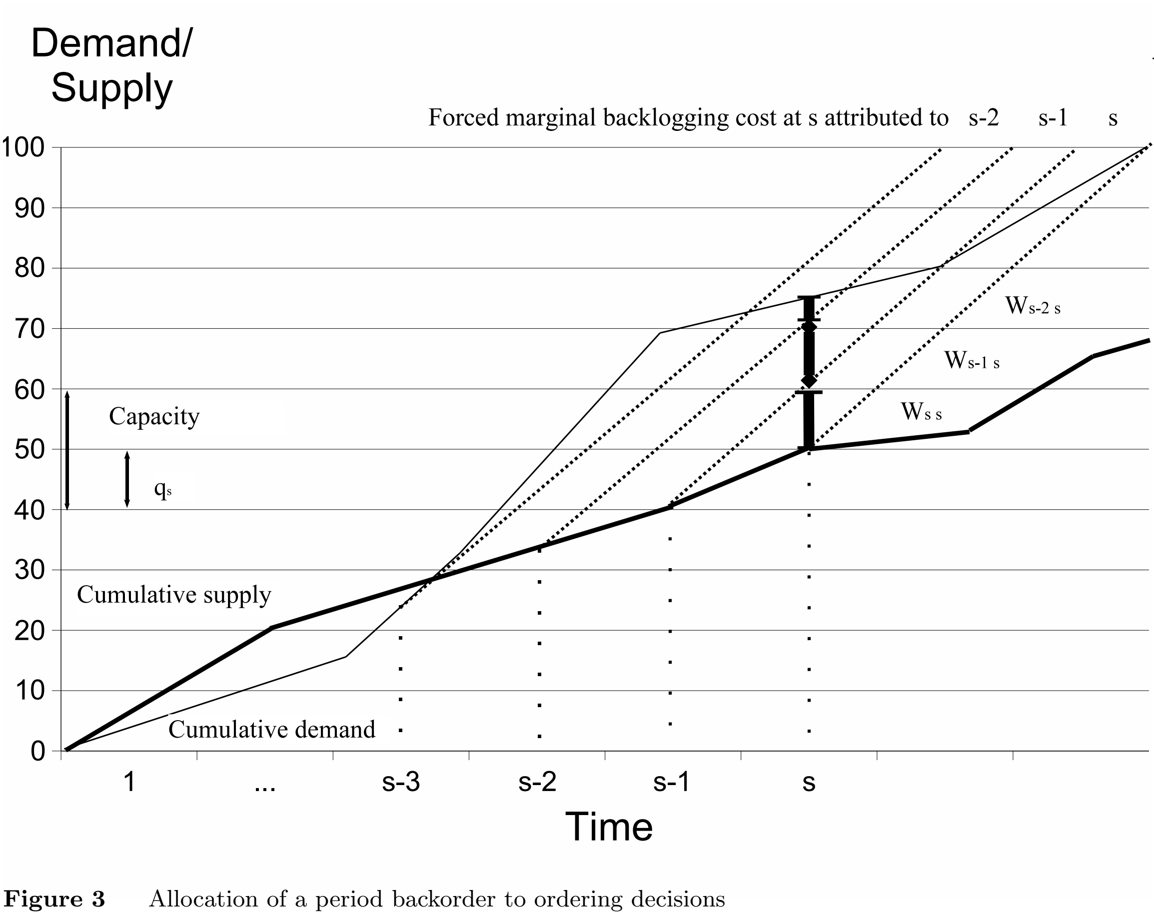

Figure 3 Allocation of a period backorder to ordering decisions

complexity of evaluating the marginal holding cost can vary depending on the level of information we assume on the demand distributions and their characteristics. In all of the common scenarios there exist straightforward methods to solve this problem efficiently (see also Hurley et al. (2006) for more details). Since in the presence of positive lead times, even computing a simple Myopic policy requires the same knowledge on the distribution of the accumulated demand over the lead time, the computational effort involved with computing the marginal holding cost is of the same order of magnitude as for the Myopic policy. Evaluating the marginal backlogging costs based on the scheme developed in this paper is analogous to the marginal holding cost. It is a sum of partial expectations of simple piecewise linear functions, and therefore, is no more difficult to compute.

Finally, observe that for uncapacitated models with $u_s = \infty$ for each s (and hence $\bar{q}_s = \infty$), our backlogging cost accounting is in fact identical to the traditional backlogging accounting discussed above. This implies that the cost accounting scheme proposed in this paper is a generalization of the one introduced in Levi et al. (2007). Therefore, the preceeding discussion is also a generalization of the corresponding algorithm and analysis in Levi et al. (2007).

4. Dual-Balancing Policy

In this section, we describe a new policy for the capacitated periodic-review stochastic inventory control problem. As in Levi et al. (2007), we call it a *Dual-Balancing policy*. We shall show that this policy has a worst-case performance guarantee of 2, i.e., for each instance of the problem, the expected cost of the policy is at most twice the expected cost of an optimal policy. Recall the assumption discussed in Section 2 that the cost parameters imply no motivation for holding inventory or backorders. This implies that, without loss of generality, for each $t = 1, \ldots, T$, $c_t = 0$ and h_t, $p_t \geq 0$. Moreover, we first describe the algorithm, its analysis, and several extensions, under the latter assumption. Then in Section 4.2 we discuss in detail the generality of this assumption.

The Dual-Balancing policy presented in this paper is based on a balancing idea similar to the one used in Levi et al. (2007) for the uncapacitated model. That Dual-Balancing policy balances, in each period s and conditioned on the observed information set f_s, the expected marginal holding cost of the units ordered in the period against the expected (traditional) backlogging cost in period $s+L$, a lead time ahead of s. However, it is readily seen that this approach does not work in the case where there is a capacity constraint on the size of the order in period s. For one, the order size q'_s that balances these two costs might not be reachable when $q'_s > u_s$.

In turn, we consider the forced marginal backlogging cost accounting and the corresponding cost it associates with period s as described in Section 3 above. Conditioned on the observed information set f_s, we now balance the expected marginal holding cost of the units ordered in period s against the expected *marginal* backlogging costs associated with period s. We will use the superscript B to refer to the Dual-Balancing policy. For each period $s = 1,\dots,T-L$, conditioning on the observed information set f_s, let $l_s^B(q_s^B)$ be the expected holding cost incurred over $[s,T]$ by the units ordered by the Dual-Balancing policy in period s. That is, $l_s^B(q_s^B) := E[H_s^B(q_s^B)|f_s]$. In Section 3 we have defined $H_s^B = \sum_{j=s+L}^{T} h_j(Q_t^B - (D_{[s,j]} - X_s^B)^+)^+$ (recall that we assume $c_s = 0$). In addition, let $\tilde{\pi}_s^B := E[\tilde{\Pi}_s^B(q_s^B)|f_s]$ be the expected backlogging cost associated with period s by the forced marginal backlogging cost accounting scheme described above, again conditioned on the observed information set f_s. Recall that in Section 3 we have defined $\tilde{\Pi}_s^B = \sum_{t=s+L}^{T} p_t W_{st}^B$ where,

$$W_{st}^B = \min\{\bar{Q}_s^B,\ (D_{[s,t]} - (X_s^B + Q_s^B + \sum_{j\in(s,t]} u_j))^+\} =$$
$$(D_{[s,t]} - (X_s^B + Q_s^B + \sum_{j\in(s,t-L]} u_j))^+ - (D_{[s,t]} - (X_s^B + \sum_{j\in[s,t-L]} u_j))^+.$$

Since if we condition on f_s, the inventory position at the beginning of period s, x_s^B, is known deterministically; it is clear that $l_s^B(q_s^B)$ and $\tilde{\pi}_s^B(q_s^B)$ are both indeed functions of q_s^B, the number of units ordered in period s.

We first discuss the case where the orders are allowed to be fractional. This implies that the functions $l_s^B(q_s^B)$ and $\tilde{\pi}_s^B(q_s^B)$ are continuous. In each period $s = 1,\dots,T-L$, given the observed information set f_s, the Dual-Balancing policy will order $q_s^B = q'_s \le u_s$ units such that the expected marginal ordering and holding cost incurred by these units over $[s,T]$ is equal to the expected forced marginal backlogging cost associated with period s. In other words, we order q'_s units such that $l_s^B(q'_s) = E[H_s^B(q'_s)|f_s] = \tilde{\pi}_s^B(q'_s) = E[\tilde{\Pi}_s^B(q'_s)|f_s]$. Next we show that this policy is well-defined. It is readily verified that $l_s^B(q_s^B)$ is a convex increasing function of q_s^B that is equal 0 for $q_s^B = 0$ and goes to ∞ as q_s^B goes to ∞. Similarly, one can verify that $\tilde{\pi}_s^B(q_s^B)$ is a decreasing convex function of q_s^B that has a non-negative value at $q_s^B = 0$ and that is equal to 0 for $q_s^B = u_s$ (in this case there is no unused slack capacity at s and $\bar{q}_s^B = 0$). Our assumption that these functions are continuous implies that q'_s, as defined above, always exists.

Computationally, q'_s is the minimizer of the function $g_s(q_s^B) := \max\{l_s^B(q_s^B), \tilde{\pi}_s^B(q_s^B)\}$, which is a convex function of q_s^B, since it is the maximum of two convex functions. Hence, in each period s, we need to solve a convex minimization problem of a single variable. In particular, if for each $j \ge s$, $D_{[s,j]}$ is distributed according to any of those distributions that are commonly used in inventory theory, then it is extremely easy to evaluate the functions $l_s^B(q_s^B)$ and $\tilde{\pi}_s^B(q_s^B)$. More generally, the complexity of the algorithm is of order T (i.e., number of time periods) times the complexity of solving the single variable convex minimization defined above. The complexity of this minimization problem can vary depending on the level of information we assume on the demand distributions and their characteristics. In all of the common scenarios there exist straightforward methods to solve this problem efficiently. In particular, q'_s is determined by the intersection of two monotone convex functions, which suggests that bisection methods can be effective in computing q'_s. We note that the Dual-Balancing policy is not a state-dependent base stock policy. However, it can be computed in an *on-line* manner, i.e., computing the policy action in period s does not require any knowledge on the future decisions to be made in the next periods. Moreover, unlike the Myopic policy, the Dual-Balancing policy does use available information about long term future demands.

4.1. Analysis

Next we show that, for each instance of the problem, the expected cost of the Dual-Balancing policy described above is at most twice the expected cost of an optimal policy. We will use the marginal cost accounting scheme described in Section 3 and amortize the period cost of the Dual-Balancing policy with the cost of the optimal policy.

Using the marginal cost accounting scheme discussed in Section 3, the expected cost of the Dual-Balancing policy can be expressed as $E[\mathcal{C}(B)] = \sum_{t=1}^{T-L} E[H_t^B + \tilde{\Pi}_t^B]$. For each $t = 1, \ldots, T-L$, let Z_t be the *random balanced cost* by the Dual-Balancing policy in period t, i.e., $Z_t = E[H_t^B | \mathcal{F}_t]$. Note that Z_t is a function of the observed information set in period t. In the next lemma we obtain an expression for the expected cost of the Dual-Balancing policy using the Z_t variables. The proof is identical to the proof of Lemma 4.1 in Levi et al. (2007).

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

In the next two lemmas we show that the cost of OPT can be amortized against some of the cost of the Dual-Balancing policy. In particular, they imply that the expected cost of OPT is at least $\sum_{t=1}^{T-L} E[Z_t]$. For each realization of the demands $D_1, \ldots, D_T$, let $\mathcal{T}_H$ be the set of periods $t = 1, \ldots, T-L$ in which the optimal policy had inventory position higher than that of the Dual-Balancing policy, i.e., the set of periods $1 \le t \le T-L$ such that $Y_t^B < Y_t^{OPT}$. Let $\mathcal{T}_\Pi$ be the set of period in which the Dual-Balancing had inventory position at least as high as that of OPT, i.e., the set of periods $t = 1, \ldots, T-L$ such that $Y_t^B \ge Y_t^{OPT}$. (We consider only the periods $t = 1, \ldots, T-L$, because the effective ordering decisions are made in these periods. Specifically, each order placed after period $T-L$ will arrive after period T.) Observe that $\mathcal{T}_H$ and $\mathcal{T}_\Pi$ are random sets that induce a random partition of the horizon.

The next lemma shows that, with probability 1, the marginal holding cost incurred by the Dual-Balancing policy in periods $t \in \mathcal{T}_H$ is at most the overall holding cost incurred by OPT, denoted by H^{OPT}, i.e., $\sum_{t \in \mathcal{T}_H} H_t^B \le H^{OPT}$ with probability 1. The proof is identical to the proof of Lemma 4.2 in Levi et al. (2007).

Lemma 2. *For each realization $f_T \in \mathcal{F}_T$, the total marginal holding cost incurred by the Dual-Balancing policy for all of the periods $t \in \mathcal{T}_H$ is at most the overall holding cost incurred by OPT, denoted by H^{OPT}, i.e., $\sum_{t \in \mathcal{T}_H} H_t^B \le H^{OPT}$ with probability 1.*

The next lemma shows that, with probability 1, the marginal backlogging cost of the Dual-Balancing policy associated with periods $t \in \mathcal{T}_\Pi$ is at most the overall backlogging penalty incurred by OPT, denoted by $\tilde{\Pi}^{OPT}$.

Lemma 3. *For each realization $f_T \in \mathcal{F}_T$, the total marginal backlogging cost of the Dual-Balancing policy associated with all of the periods $t \in \mathcal{T}_\Pi$ is at most the overall backlogging penalty incurred by OPT, denoted by $\tilde{\Pi}^{OPT}$, i.e., $\sum_{t \in \mathcal{T}_\Pi} \tilde{\Pi}_t^B \le \tilde{\Pi}^{OPT}$ with probability 1.*

T he forced marginal backlogging cost associated with the periods in $\mathcal{T}_\Pi$ is equal to

$$\sum_{s \in \mathcal{T}_\Pi} \sum_{t: t \ge s+L} p_t W_{st}^B = \sum_t p_t \sum_{s \in \mathcal{T}_\Pi : s \le t-L} W_{st}^B.$$

Therefore, it is sufficient to show that for each $t=L+1,\dots,T$, the traditional backlogging cost incurred by OPT in that period is at least as much as the forced backlogging costs incurred by the Dual-Balancing policy in period t as a result of decisions made in periods $\{s\in\mathcal{T}_\Pi : s\le t-L\}$. In other words, it is sufficient to show that for each $t=L+1,\dots,T$, we have

$$(D_t - NI_t^{OPT})^+ \ge \sum_{s\in\mathcal{T}_\Pi : s\le t-L} W_{st}^B,$$

with probability 1. (Recall that the backlogging costs over the periods $1,\dots,L$ are the same for all policies.)

Consider now a specific realization $f_T\in\mathcal{F}_T$ and some period $t=1,\dots,T$. If there is no period in $\{s\in\mathcal{T}_\Pi : s\le t-L\}$ with $w_{st}^B>0$, then there is nothing to prove. Assume that such a period s exists, and let s_l and s_e be, respectively, the latest and the earliest periods in the set $\{s\in\mathcal{T}_\Pi : s\le t-L,\ w_{st}^B>0\}$, respectively (it is possible that $s_l=s_e$). We note again that here we abuse our notation and consider the set $\mathcal{T}_\Pi$ as the realized set of periods according to the specific realization f_T. In particular, s_e and s_l are the respective realizations of random variables S_e and S_l. We have already seen (in the discussion in Section 3) that for each $s\in(s_e,s_l]$ we have $w_{st}^B=\bar{q}_s^B$, and $w_{s_e,t}^B\le d_{[s_e,t]}-(x_{s_e}+q_{s_e}^B+\sum_{j\in(s_e,t-L]}u_j)$. Indeed,

$$\begin{aligned}
d_t - ni_t^{OPT} &= d_t - (y_{s_l}^{OPT} + \sum_{j\in(s_l,t-L]} q_j^{OPT} - d_{[s_l,t)}) \ge d_{[s_l,t]} - (y_{s_l}^B + \sum_{j\in(s_l,t-L]} u_j) \\
&= d_{[s_l,t]} - (y_{s_e}^B + \sum_{j\in(s_e,s_l]} q_j^B - d_{[s_e,s_l)} + \sum_{j\in(s_l,t-L]} u_j) \\
&= d_{[s_e,t]} - (x_{s_e}^B + q_{s_e}^B + \sum_{(s_e,t-L]} u_j) + \sum_{j\in(s_e,s_l]} \bar{q}_j^B \\
&\ge \sum_{j\in[s_e,s_l]} w_{st}^B \ge \sum_{j\in[s_e,s_l]\cap\mathcal{T}_\Pi} w_{st}^B.
\end{aligned}$$

The first equality is based again on the fact that for each feasible policy and for each $s\le t$, we have $NI_t=Y_s+\sum_{j\in(s,t-L]}Q_j-D_{[s,t)}$, applied to OPT and periods $s_l\le t-L$. The first inequality follows from the assumption that $s_l\in\mathcal{T}_\Pi$ and so $y_{s_l}^{OPT}\le y_{s_l}^B$, and from the capacity constraints that imply $q_j^{OPT}\le u_j$. The second equality follows from the fact that (for each $s\le s'$) $Y_{s'}=Y_s+\sum_{j\in(s,s']}Q_j-D_{[s,s')}$ applied to the Dual-Balancing policy and periods $s_e\le s_l$. The last equality is achieved by adding and subtracting $\sum_{j\in(s_e,s_l]}\bar{q}_j^B$ and from the fact that $u_j=Q_j+\bar{Q}_j$. The proof then follows.

As a corollary of Lemmas 1, 2 and 3 we get the following theorem.

THEOREM 2. *The Dual-Balancing policy has a worst-case performance guarantee of 2, i.e., for each instance of the capacitated periodic-review stochastic inventory control problem, the expected cost of the Dual-Balancing policy is at most twice the expected cost of an optimal solution, i.e.,* $E[\mathcal{C}(B)]\le 2E[\mathcal{C}(OPT)]$.

F rom Lemma 1, we know that the expected cost of the Dual-Balancing policy is equal to twice the expected cost of the sum of the Z_t variables, i.e., $E[\mathcal{C}(B)]=\sum_{t=1}^{T-L}E[Z_t]$. From Lemmas 2 and 3 we know that, with probability 1, the cost of OPT is at least as much as the holding cost incurred by units ordered by the Dual-Balancing policy in periods $t\in\mathcal{T}_H$ plus the forced marginal backlogging cost of the Dual-Balancing policy that is associated with periods $t\in\mathcal{T}_\Pi$. In other words, with probability 1, $H^{OPT}+\tilde{\Pi}^{OPT}\ge\sum_{t\in\mathcal{T}_H}H_t^B+\sum_{t\in\mathcal{T}_\Pi}\tilde{\Pi}_t^B$. Using again conditional expectations and the definition of Z_t, this implies that indeed,

$$E[\mathcal{C}(OPT)] \geq E[\sum_{t \in \mathcal{T}_H} H_t^B + \sum_{t \in \mathcal{T}_\Pi} \tilde{\Pi}_t^B] = \sum_t E[H_t^B \cdot \mathbb{1}(t \in \mathcal{T}_H) + \tilde{\Pi}_t^B \cdot \mathbb{1}(t \in \mathcal{T}_\Pi)] =$$
$$\sum_t E[E[H_t^B \cdot \mathbb{1}(t \in \mathcal{T}_H) + \tilde{\Pi}_t^B \cdot \mathbb{1}(t \in \mathcal{T}_\Pi) | \mathcal{F}_t]] = \sum_t E[(\mathbb{1}(t \in \mathcal{T}_H) + \mathbb{1}(t \in \mathcal{T}_\Pi)) Z_t] = \sum_t E[Z_t].$$

We note that if the optimal policy is deterministic (i.e., it makes deterministic decisions in each period t given the observed information set f_t), then if we condition on $\mathcal{F}_t$, then y_t^B and y_t^{OPT} are known deterministically, and so are the indicators $\mathbb{1}(t \in \mathcal{T}_H)$ and $\mathbb{1}(t \in \mathcal{T}_\Pi)$. If the optimal policy is random, then the same arguments above still work. We now need to condition not only on $\mathcal{F}_t$ but also on the decisions made by the policies. Since the inventory control policy does not have any effect on the evolution of the demand, the arguments above are still valid. This concludes the proof of the theorem.

We note that the examples discussed in Levi et al. (2007) show that the above analysis is tight. However, the analysis hints that in a typical scenario, the performance would be significantly better. Hurley et al. (2006) present a thorough empirical analysis of the typical performance of Dual-Balancing policies in uncapacitated models. In Section 6, we present empirical results that confirm that this phenomenon extends to the capacitated case.

Finally, we note that the Dual-Balancing policies and the worst-case analysis can be extended to models where the capacities in each period are generated by some exogenous random process, and the exact capacity available in period t is observed only at the beginning of the period. Thus, the Dual-Balancing policies provide a worst-case guarantee of 2 for this important extension as well. In this case, the expectations of the marginal backlogging costs are taken with respect to both the random future demands and random future capacities. In Appendix C, we consider two extensions of the Dual-Balancing policy and the worst-case analysis. Specifically, we discuss the extensions to models where orders must be integral and the demands are integer-valued random variables, and to models with stochastic lead times under the no order crossing assumption.

4.2. Cost Transformation

In this section, we discuss in detail the cost transformation that enables us to assume, without loss of generality, that for each period $t = 1, \ldots, T$, we have $c_t = 0$ and $h_t,\ p_t \geq 0$. Consider any instance of the problem with cost parameters that imply no speculative motivation for holding inventory or backorders (as discussed in Section 2). Following Levi et al. (2007), we use a simple transformation of the cost parameters to construct an equivalent instance, with the property that for each period $t = 1, \ldots, T$, we have $c_t = 0$ and $h_t,\ p_t \geq 0$. More specifically, the modified instance has the same set of optimal policies. Applying the Dual-Balancing policy to that instance, we obtain a policy that is different from the original dual balancing policy, and which also has a performance guarantee of at most 2 with respect to the original problem. We shall show that this cost transformation can improve the performance guarantee of the Dual-Balancing policy in cases where the ordering cost is the dominant part of the overall cost. In practice this is often the case.

We now describe the transformation for the case with no lead time ($L = 0$) and $\alpha = 1$; the extension to the case of arbitrary lead time is straightforward. Recall that any feasible policy P satisfies, for each $t = 1, \ldots, T$, $Q_t = NI_t - NI_{t-1} + D_t$ (for ease of notation we omit the superscript P). Using these equations, we can express the ordering cost in each period t as $c_t(NI_t - NI_{t-1} + D_t)$. Now replace NI_t with $NI_t^+ - NI_t^-$, its respective positive and negative parts.

This leads to the following transformation of cost parameters. We let $\hat{c}_t := 0,\ \hat{h}_t := h_t + c_t - c_{t+1}$ ($c_{T+1} = 0$) and $\hat{p}_t := p_t - c_t + c_{t+1}$. Note that the assumptions on the cost parameters $c_t,\ h_t$, and p_t discussed in Section 2, and in particular, the assumption that there is no speculative motivation to hold inventory or backorders, imply that $\hat{h}_t$ and $\hat{b}_t$ above are non-negative ($t = 1, \ldots, T$). Observe that the parameters $\hat{h}_t$ and $\hat{b}_t$ will still be non-negative even if the parameters c_t, h_t, and p_t are negative and as long as the above assumption holds. Moreover, this enables us to incorporate into the model a negative salvage cost at the end of the planning horizon (after the cost transformation we will have non-negative cost parameters). It is readily verified that the induced problem is equivalent to the original one. More specifically, for each realization of the demands, the cost of each feasible policy P in the modified input decreases by exactly $\sum_{t=1}^{T} c_t d_t$ (compared to its cost in the original input). Therefore, any optimal policy for the modified input is also optimal for the original input.

Now apply the Dual-Balancing policy to the modified problem. We have seen that the assumptions on c_t, h_t and p_t ensure that $\hat{h}_t$ and $\hat{p}_t$ are non-negative and hence the analysis presented above is valid. Let opt and $\overline{opt}$ be the optimal expected cost of the original and modified inputs, respectively. Clearly, $opt = \overline{opt} + E[\sum_{t=1}^{T} c_t D_t]$. Now the expected cost of the Dual-Balancing policy in the modified input is at most $2\overline{opt}$. Its cost in the original input is then at most $2\overline{opt} + E[\sum_{t=1}^{T} c_t D_t] = 2opt - E[\sum_{t=1}^{T} c_t D_t]$. This implies that if $E[\sum_{t=1}^{T} c_t D_t]$ is a large fraction of opt, then the performance guarantee of the expected cost of the Dual-Balancing policy might be significantly better than 2. For example, if $E[\sum_{t=1}^{T} c_t D_t] \geq 0.5opt$, then we can conclude that the expected cost of the Dual-Balancing policy is at most $1.5opt$. It is indeed the case in many real life problems that a major fraction of the total cost is due to the ordering cost. The intuition of the above transformation is that $\sum_{t=1}^{T} c_t D_t$ is a cost that any feasible policy must pay. As a result, we treat it as an invariant in the cost of any policy and apply the approximation algorithm to the rest of the cost.

In the case where we have a lead time L, we use the equations $Q_t := NI_{t+L} - NI_{t+L-1} + D_{t+L}$, for each $t = 1, \ldots, T - L$, to get the same cost transformation. The transformation for $\alpha > 1$ is also straightforward. Also, it is not hard to see that the cost transformation can be modified to remove, say, $\gamma\%$ of the per-unit ordering costs, where $0 < \gamma < 100$. This leads to a continuum of dual balancing policies, all of which are 2-approximations.

5. Improved Policy & Bounds on the Optimal Inventory Levels

In this section, we consider two semi-myopic (modified) base-stock policies that are easy to compute in an on-line manner and provide, respectively, lower bounds and upper bounds on the inventory levels of an optimal policy y_t^{OPT}, in each period $t = 1, \ldots, T$. We believe that these bounds can be used effectively to improve existing algorithms for computing inventory control policies for the capacitated model discussed in this paper and other capacitated stochastic inventory models. Moreover, as in Hurley et al. (2006), we shall show that these policies provide bounds that are strong in the following sense: each policy that, for some period t and some state f_t, has inventory level outside the range defined by the respective lower and upper bounds can be improved. In particular, there is another (modified) policy that in period t and state f_t, admits an inventory level within the specified range, with expected cost no greater than the expected cost of the original policy. In other words, any policy that violates these respective bounds is dominated by another policy. We then follow Hurley et al. (2006) and construct an *Improved Dual-Balancing policy* that incorporates these bounds. This policy also has a performance guarantee of 2 and as the computational study for the uncapacitated model in Hurley et al. (2006) suggests, we expect that it will have a better typical performance.

The policies we consider are called *Lower-Myopic* (denoted by LM) and *Upper-Myopic* (denoted by UM), respectively. In the Lower-Myopic policy, in each period s, conditioning on the observed information set f_s, we minimize the *sum* of the expected marginal holding cost of the units ordered in that period and the traditional expected backlogging costs a lead time ahead. That is, in each period s, we minimize

$$g_s^{LM}(q_s) = l_s^{LM}(q_s) + E[p_{s+L}(D_{[s,s+L]} - (x_s + q_s))^+ | f_s],$$

under the constraint $q_s \le u_s$. This is a convex function of q_s. This policy has been first proposed for the uncapacitated model by Levi et al. (2007) who called it the *Minimizing policy*. They have shown that this is a base-stock policy that provides lower bounds on the optimal base-stock levels. However, in the capacitated model it is possible that the actual minimizer will not be attainable. In this case we order up to capacity, and this provides a modified base-stock policy. In this paper, we extend and generalize their proof for the capacitated model. In the Upper-Myopic policy, in each period s, again conditioning on f_s, we minimize the sum of the expected period holding cost and the expected forced marginal backlogging. Thus, we minimize

$$g_s^{UM}(q_s) = \tilde{\pi}_s^{UM}(q_s) + E[h_{s+L}(x_s + q_s - D_{[s,s+L]})^+ | f_s],$$

subject to $0 \le q_s \le u_s$, which is also convex in q_s. We shall show that this policy provides upper bounds on the inventory levels of an optimal policy. By arguments similar to the ones used by Levi et al. (2007), it can be shown that this gives rise to yet another modified base-stock policy. (In particular, $g_s^{UM}(q^1) - g_s^{UM}(q^2)$ depends only on $y^1 = x_s + q^1$ and $y^2 = x_s + q^2$.) To the best of our knowledge, this is a new way for deriving upper bounds on the inventory levels of an optimal policy in the capacitated model. We note that it is not clear whether the classical Myopic policy, where we minimize the expected period cost, provides any bounds for capacitated models. Another similar open question is how the policy that in each period minimizes the sum of the expected marginal holding cost and expected forced marginal backlogging cost is related to an optimal policy.

Let Y_t^{LM} and Y_t^{UM} be the respective inventory position (after orders are placed) of the Lower-Myopic and the Upper-Myopic policies in period $t = 1, \ldots, T$. Specifically, we assume that Y_t^{LM} is the smallest minimizer of the corresponding period problem being solved (see above) and that Y_t^{UM} is the largest minimizer of the corresponding period problem. Note that the inventory position levels depend on the specific state (f_t, x_t), but for ease of notation we omit the indication of the state. The two semi-myopic policies described above can be implemented in an on-line manner, i.e., regardless of the action control in future periods. We shall show that for each evolution f_T, these two policies provide lower and upper bounds on the inventory levels of any optimal policy, i.e., $Y_t^{LM} \le Y_t^{OPT} \le Y_t^{UM}$, with probability 1, for each $t = 1, \ldots, T$. Moreover, we shall show that each non-dominated policy P must have $Y_t^{LM} \le Y_t^P \le Y_t^{UM}$, for each $t = 1, \ldots, T$.

The next two lemmas show that each policy P that has, for some period s and state f_s, inventory position $y_s^P \notin [y_s^{LM}, y_s^{UM}]$, can be strictly improved by a modified policy P' with $y_s^{P'} \in [y_s^{LM}, y_s^{UM}]$ and expected cost at most the expected cost of P. For the sake of simplicity, we consider a model with no lead time (the extensions to the case with $L > 0$ are straightforward).

LEMMA 4. *Consider a feasible policy P, and suppose that for some period s and information set f_s, we have $y_s^P < y_s^{LM}$. Further assume that s is the earliest such period. Then the policy P' that follows P until period $s-1$, then orders up to y_s^{LM} in period s and again imitates P over the interval $(s, T]$, has expected cost no larger than the expected cost of P.*

S ince P' follows P over $[1, s)$, we conclude that they incur exactly the same cost over that interval, and that they have the same inventory position $x_s \le y_s^P < y_s^{LM}$. Since s is the first such period, we conclude that P' can indeed order up to y_s^{LM}. Now over $(s, T]$, P' imitates P; that is, it orders nothing if $X_j^{P'} \ge Y_j^P$ and orders up to Y_j^P otherwise (for each $j \in (s, T]$). Moreover, the policy P' has ordered $q_s^{P'}$ units in period s. Consider the overall expected marginal holding cost of these units and the expected (traditional) backlogging cost incurred by P' in period s. By the definition of $q_s^{P'}$, it is clear that this is no greater than the expected marginal holding cost and expected (traditional) backlogging cost incurred by the policy P in period s. For each period $j \in (s, T]$, we know that with probability 1, $Y_j^{P'} \ge Y_j^P$ and that $Q_j^{P'} \le Q_j^P$. This implies that the backlogging incurred by policy P' over that interval is no greater than the backlogging cost incurred by policy P, and similarly, the marginal holding cost policy P' incurs over that interval is no greater than the respective marginal holding cost of policy P. The lemma then follows.

LEMMA 5. *Consider a feasible policy P, and suppose that for some period s and information set f_s, we have $y_s^P > y_s^{UM}$. Further assume that s is the earliest such period. Then the policy P' that follows P until period $s-1$, then orders up to y_s^{UM} in period s and again imitates P over the interval $(s,T]$, has expected cost no larger than the expected cost of P.*

B y arguments identical to the ones in Lemma 4, we conclude that P' and P incur the same cost over $[1,s)$ and that they have the same inventory position $x_s \le y_s^{UM} < y_s^P$. The first inequality follows from the fact that s is the first period in which P has more inventory than the Upper-Myopic policy. Thus, P' can order up to Y_s^{UM}, and assume that it orders $q_s^{p'}$. Consider the overall expected forced marginal backlogging cost and expected period holding cost incurred in period s by policy P'. By the definition of $q_s^{P'}$, we conclude that this expected cost is smaller than the respective expected cost incurred by policy P in period s. Now over $(s,T]$ P' again tries to imitate P, i.e., for each $j \in (s,T]$, it will order up to Y_j^P or up to the capacity u_j. Now let S' be the earliest (random) period after period s in which P' has reached $Y_{S'}^P$. Clearly, over $(S',T]$ the policies P' and P are again identical and hence, incur the same cost. Observe that, for each $j \in (s,S']$, we have $Y_j^{P'} \le Y_j^P$ and $\bar{Q}_j^{P'} \le \bar{Q}_j^P$, with probability 1. This implies that the expected holding cost and the expected forced marginal backlogging penalty incurred by policy P' over that interval are each no greater than the respective expected cost incurred by policy P. The lemma then follows.

Lemmas 4 and 5 imply the following corollary.

COROLLARY 1. *For any optimal policy and for each complete evolution f_T, the Lower-Myopic and Upper-Myopic policies provide respective lower and upper bounds on the inventory levels of the optimal policy, i.e., $Y_t^{LM} \le Y_t^{OPT} \le Y_t^{UM}$ with probability 1, for each $t = 1,\ldots,T$.*

Now consider the Improved Dual-Balancing policy denoted by superscript IB. In each period s, given the observed information set f_s and the inventory position at the beginning of the period, we still consider balancing the expected marginal holding cost against the expected marginal backlogging cost, and compute q_s' as described in Section 4. (That is, given the observed information set f_s and the inventory position at the beginning of period s, ordering q_s' will balance the expected marginal holding cost and the expected marginal forced backlogging costs associated with period s.) However, in each case where the original balancing quantity brings the inventory position below y_s^{LM} (i.e., $x_s^{IB} + q_s' < y_s^{LM}$) or above y_s^{UM} (i.e., $x_s^{IB} + q_s' > y_s^{UM}$), we fix this decision by instead increasing the order up to y_s^{LM} or decreasing it down to y_s^{UM}, respectively. It can be readily verified that for each evolution f_T and each period s, we have $y_s^{LM} \le y_s^{IB} \le y_s^{UM}$.

We next prove the following theorem.

THEOREM 3. *The Improved Dual-Balancing policy has a performance guarantee of 2.*

O bserve that in the Improved Dual-Balancing policy it is no longer true that, in each period t, the expected marginal holding cost is equal to the expected forced marginal backlogging cost. Now let Z_t be the maximum among the expected marginal holding cost and expected forced marginal backlogging cost, i.e., $Z_t = \max\{E[H_t^{IB}(Q_t^{IB})|\mathcal{F}_t], E[\tilde{\Pi}_t^{IB}(Q_t^{IB})|\mathcal{F}_t]\}$, where Q_t^{IB} is the order quantity placed by the Improved Dual-Balancing policy in period s. (As already mentioned Q_t^{IB} can be either larger or smaller than the balancing quantity Q_t'.) Similar to Lemma 1, we now conclude that $E[\mathcal{C}(IB)] \le 2\sum_t E[Z_t]$.

Next we modify the definition of the sets $\mathcal{T}_H$ and $\mathcal{T}_\Pi$ in Section 4. The set $\mathcal{T}_H$ will consist of periods $t = 1,\ldots,T-L$ such that (i) $Y_t^{LM} < Y_t^{IB} < Y_t^{UM}$ and $Y_t^{IB} \le Y_t^{OPT}$; or (ii) $Y_t^{IB} = Y_t^{LM} < Y_t^{UM}$; or (iii) $Y_t^{IB} = Y_t^{LM} = Y_t^{UM} = Y_t^{OPT}$ and the Improved Dual-Balancing policy orders more than the balancing quantity Q_t'. (That is, $X_t^{IB} + Q_t' \le Y_t^{LM} = Y_t^{UM}$ and $Q_t^{IB} \ge Q_t'$.) The set $\mathcal{T}_\Pi$ will consist of all the other periods in $t = 1,\ldots,T-L$. Specifically, $\mathcal{T}_\Pi$ contains periods such that (i) $Y_t^{LM} < Y_t^{IB} < Y_t^{UM}$ and $Y_t^{IB} > Y_t^{OPT}$; or (ii) $Y_t^{LM} < Y_t^{IB} = Y_t^{UM}$; or (iii) $Y_t^{IB} = Y_t^{LM} = Y_t^{UM} = Y_t^{OPT}$ and the Improved Dual-Balancing policy orders less than the balancing quantity Q_t'. (That is, $X_t^{IB} + Q_t' > Y_t^{UM} = Y_t^{LM}$ and $Q_t^{IB} < Q_t'$.) Note that for each $t \in \mathcal{T}_H$, we have $Y_t^{IB} \le Y_t^{OPT}$ and for each $t \in \mathcal{T}_\Pi$, we have $Y_t^{IB} \ge Y_t^{OPT}$. Thus, the arguments used to prove Lemmas 2 and 3 are still valid. It is then sufficient to show that, for each $t \in \mathcal{T}_H$, we have $E[H_t^{IB}(Q_t^{IB})|\mathcal{F}_t] = Z_t$, and, for each $t \in \mathcal{T}_\Pi$, we have $E[\tilde{\Pi}_t^{IB}(Q_t^{IB})|\mathcal{F}_t] = Z_t$. This will imply that the arguments in the proof of Theorem 2 are still valid and the performance guarantee of the policy then follows.

Assume now that for some $t \in \mathcal{T}_H$ and some $f_t \in \mathcal{F}_t$, we have $E[H_t^{IB}(q_t^{IB})|f_t] < z_t$. However, this can happen only if in that period the Improved Dual-Balancing policy orders below Q_t' and $Y_t^{IB} = Y_t^{UM}$. (The Improved Dual-Balancing policy orders $Q_t^{IB} < Q_t'$ only when $X_t^{IB} + Q_t' > Y_t^{UM}$, and then it decreases the order until $Y_t^{IB} = Y_t^{UM}$.) This leads to a contradiction since by definition $t \in \mathcal{T}_\Pi$ (see cases (ii) and (iii) in the definition of $\mathcal{T}_\Pi$ above).

Similarly, assume that for $t \in \mathcal{T}_\Pi$ and some $f_t \in \mathcal{F}_t$, we have $E[\tilde{\Pi}_t^{IB}(Q_t^{IB})|\mathcal{F}_t] < z_t$. This can happen only if in that period the Improved Dual-Balancing policy orders $Q_t^{IB} > Q_t'$ (i.e., $X_t^{IB} + Q_t' < Y_t^{LM}$) and $Y_t^{IB} = Y_t^{LM}$. However, again we get a contradiction since by definition $t \in \mathcal{T}_H$ (see cases (ii) and (iii) in the definition of $\mathcal{T}_H$ above). This concludes the proof of the lemma.

6. Computational Experiments

As we mentioned in the introduction, due to state space explosion, the corresponding inventory control models are very difficult from a computational perspective. Consequently, we study the typical performance of the balancing policies in two settings. In the first setting the optimal solution of the capacitated inventory system is easily computed, but there is no evolution of forecasts (i.e., demands are independent over time). This enables us to see how close to optimal the Balancing policy is, in at least one setting. The second experiment is more realistic, in that the demand and forecast evolution processes are governed by the multiplicative MMFE model. In fact these are the settings, in which balancing policies are most attractive, because optimal policies are inaccessible and no provably good heuristics or even reasonable lower bounds are available. As a result, we benchmark the performance of the balancing policies using the Myopic and the other semi-myopic policies developed in this paper in Section 5. In these experiments the balancing policies were very robust. For the model with independent demands, the Dual-Balancing policy came within 11% of the optimal cost on average, within 17% of optimal in 95% of the trials and never exceeded the optimal cost by more than 25%. Moreover, the balancing policies out-perform the myopic policy by 49% in the first experiment and by 27% in the second, on average. (In many scenarios the balancing policies improve upon myopic by more than 50%.) This indicate that the typical performance of the balancing policies is significantly better than the worst-case guarantees.

6.1. Experiments with Translated-Mass Exponential Demand Distributions

In this experiment we consider infinite-horizon problems with i.i.d. demand, i.e., the distribution of $(D_t|\mathcal{F}_s)$ is independent of both $\mathcal{F}_s$ and t. We assume that D_t has a translated-mass exponential distribution, meaning that $P(D_t > x) = 1$ if $x < a$, and otherwise, $P(D_t > x) = q\, e^{-\theta(x-a)^+}$, where $0 \le q \le 1$, $\theta > 0$, $a \ge 0$, and $a \cdot (1-q) = 0$. If $q = 1$ then D_t has an exponential distribution, translated to the right by a units. If $q < 1$ then $a = 0$, $D_t = 0$ with probability $1-q$, and with probability q, D_t follows an exponential distribution. For every positive mean and variance there is a unique translated-mass exponential distribution.

For infinite-horizon problems with translated-mass exponential demand, a stationary order-up-to policy is optimal. The optimal policy and its cost are easily obtained, using the following observation: for translated-mass exponential demand, the lower and upper bounds in Theorem 2 of reference Glasserman (1997) coincide.

The demand D_t has mean 1. We start with a Base Case, in which D_t has variance 1, the capacity is 1.5, and the backorder cost per day is 8 times larger than the holding cost. Figure 4 illustrates what happens when we fix two of these parameters and vary the third one. On the vertical axis we

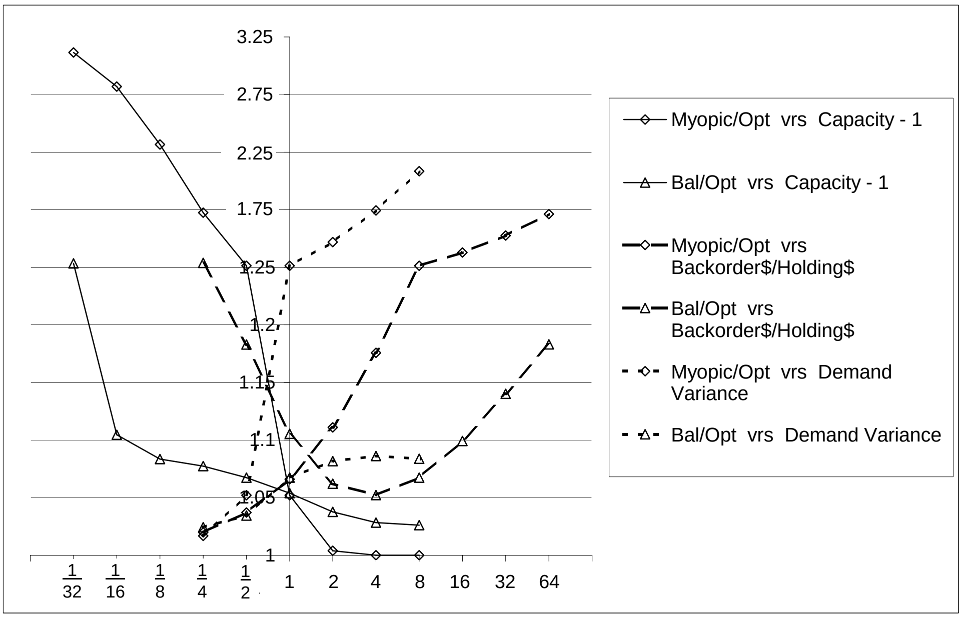

Figure 4 Sensitivity of Performance to Capacity, Backorder Costs, and Demand Variance

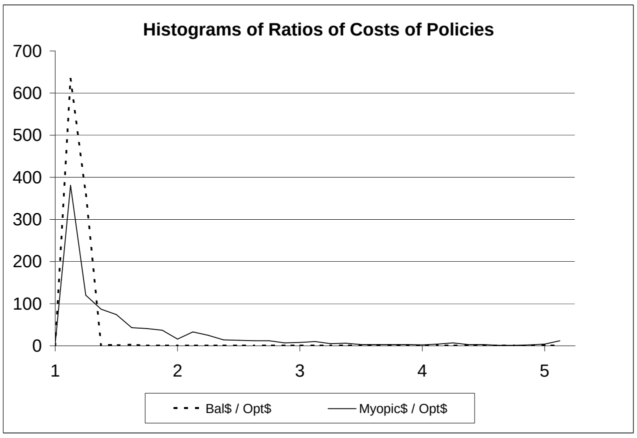

Figure 5 Histogram of Cost, as a Fraction of Optimal Cost

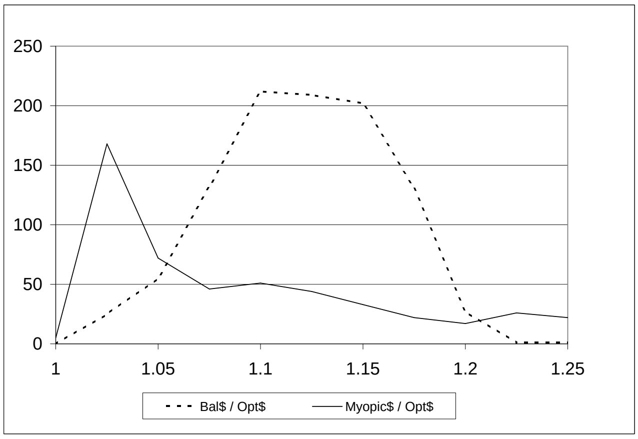

Figure 6 Detailed Histogram of Cost, as a Fraction of Optimal Cost

show the ratio of the cost of the Balancing policy to the optimal cost, and the ratio of the Myopic's cost to the optimal cost. Note that the scale on the vertical axis is not uniform. For the solid lines, the horizontal axis displays the excess capacity (i.e., the capacity minus the mean demand, or "Capacity - 1"). For the dashed lines, the horizontal ordinate is the ratio of the backorder cost per day to the holding cost. For the dotted lines, the horizontal ordinate is the variance of the demand.

In addition, we randomly generated 1000 problem instances, using a mean demand of 1. The capacity, the backorder-to-holding-cost ratio, and the standard deviation of the demand are all randomly generated from translated beta distributions. For the capacity, the distribution has minimum, maximum, mean and standard deviation equal to (1.05, 3.3, 1.61, 0.32). For the backorder-to-holding-cost ratio and the standard deviation of the demand, the corresponding values are (1, 101, 26.00, 14.43) and (0.1, 3.6, 0.98, 0.51). The computations were done using JAVA on a standard PC, and computing the balancing decision in each period took 0.00015 seconds on average.

Figure 5 shows histograms of the ratio of the balancing policy's cost and the optimal cost, and the ratio of the myopic policy's cost and the optimal cost. Figure 6 is a restricted view of Figure 5, with a finer grid, limited to the neighborhood around 1. The ratio of the Balancing policy's cost to the optimal cost is 1.11 on average, with a standard deviation of 0.049, a 95-th percentile of 1.17, and a maximum of 1.58. For the Myopic, the corresponding ratio has mean 1.60, standard deviation 0.92, 95-th percentile 3.38, and maximum 8.61. This indicates that the Balancing policy is very robust compared to the Myopic.

6.2. Experiments with Multiplicative-MMFE-based Demand and Forecast Evolution

This test uses the experimental design of Hurley et al. (2006), in which the uncapacitated version of the balancing algorithm was tested. In all of our experiments, the holding and backorder costs are $h_t = 1$ and $p_t = 10$. A horizon of length $T = 40$ was used, and forecasts of demand evolve according to the multiplicative MMFE model Heath and Jackson (1994). The mean demand per period, averaged over the 40 periods in the time horizon, is 400 in all cases. The capacity is 460 units per period.

The experimental design consists of 82 scenarios. For each scenario we tested 1000 random problem instances. The scenarios were designed to capture a variety of settings and characteristics. Demand and forecast variability can be either high or low, and lead times can be short or long. Some scenarios study different types of seasonality in the demand. Others consider product launches and product phaseouts. Some scenarios account for the fact that many forecasting systems generate accurate forecasts that extend many time periods into the future, whereas other systems can only forecast accurately in the near term. In addition, shifts in forecasts can demonstrate either no correlation, positive correlation or negative correlation. The scenarios are described in detail in Appendix B, and in Hurley et al. (2006).

We study five policies: Myopic, Lower Myopic, Upper Myopic, Improved Balancing and Improved Balancing. For each of the 82 scenarios constructed and for each policy, we examine the average per period cost of the policy over 1000 runs. Note that since we consider a complex environment and relatively long horizon ($T = 40$), it is not possible to compute the optimal expected cost. Moreover, to the best of our knowledge, it is now even known how to compute reasonable lower bounds in this setting. Instead, we use as our bench mark the Myopic policy and the other semi myopic policies discussed in Section 5. The policies were computed using MATLAB on a standard PC. The average times to compute the period ordering decisions were 0.0031, 0.0738, 0.0412 seconds for the Myopic, the Minimizing and Balancing policies, respectively.

In Figure 7 we provide histograms of the ratio of the cost of each policy, divided by the cost of Myopic. Both Balancing and Improved Balancing outperform Myopic in every one of the 82 scenarios. Relative to Myopic, they provide an average saving of 27.2% and 32.4%, respectively. Lower Myopic is very close to Myopic (ratio is usually close to 1), and is sometimes worse than Myopic. The trend is not unexpected since Myopic often under-orders in capacitated systems, and Lower Myopic always orders less than Myopic. Upper Myopic is virtually identical to Improved Balancing, which truncates the Balancing order quantities using the order-up-to levels of Upper and Lower Myopic.

In all of our computational experiments, the performance of the Balancing policy is both strong and consistent. Improved balancing is better than Balancing.

Appendix A: Marginal Backlogging Cost Accounting Approach - Numerical Example

To provide more intuition, we illustrate the new backlogging cost accounting through a simple example. Suppose that the order capacity is 5 in all periods, $L = 0$ and $\alpha = 1$. Assume that the

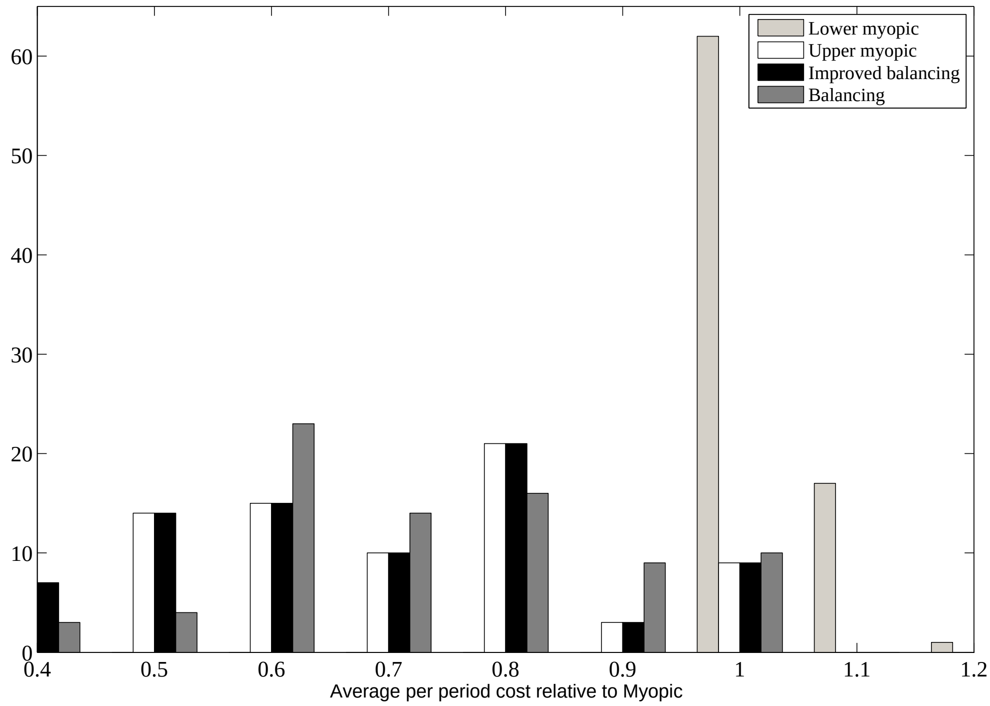

Figure 7 Performance of Four Policies relative to Myopic under Forecast Evolution

inventory position at the beginning of period 3 was $x_3 = 3$, and that we have ordered $q_3 = 3$, $q_4 = 5$, $q_5 = 4$ and $q_6 = 2$ units in periods 3, 4, 5 and 6, respectively. Now say that the demands were $d_3 = 3$, $d_4 = 3$, $d_5 = 5$ and $d_6 = 11$ in periods 3, 4, 5 and 6, respectively. In particular, the accumulated demand over periods $[3,6]$, $d_{[3,6]}$, is equal to 22. This implies that in period 6 we had a shortage of 5 units, each of which incurred a penalty cost of p_t at the end of period 6. Out of these 5 units of shortage at the end of period 6, we associate a backlogging penalty of 3 units of shortage with period 6 (the unused slack capacity in this period is 3), a penalty of 1 unit of shortage with period 5 (the unused slack capacity in this period is 1), no cost is associated with period 4 since we ordered up to capacity, and finally the penalty of 1 units of shortage is associated with period 3 ($d_{[3,6]} - (3+3+5+5+5) = 1$). In other words, $w_{36} = 1$, $w_{46} = 0$, $w_{56} = 1$ and $w_{66} = 3$. This example illustrates how we backtrack the 'source' of each unit of shortage and its corresponding backlogging cost incurred in period t, and associate it as forced backlogging cost to past periods. If $L > 0$, then we start the backtracking in period $t - L$, since only orders in periods earlier than $t - L + 1$ could have arrived by time t.

Appendix B: Experimental Design

In this appendix we give a detailed description of the scenarios that form the basis of the experiments done in Section 6.2. The space of potential parameter settings for this study is very large. In addition to parameters describing the inventory system, there are many parameters that describe the manner in which forecasts of demand evolve over time. A fully comprehensive study is beyond the scope of this paper. Our goal is to study a broad range of potential application settings, with emphasis on the demand and forecasting processes.

The experimental design is oriented around a Base Case and six sets of *scenarios*, each of which expands the Base Case in an interesting dimension. In each set of scenarios we vary specific input parameters. The first three of these scenario sets study *first-order effects*, in this case, trends and seasonality patterns in the demand. The final three scenario sets study *second order effects* by varying the probability model that governs the variance of the demand and of the forecast errors, and the correlations that exist between them.

We begin by reviewing the structure, and some of the notation, of the MMFE model. Then we discuss the Base Case. After that we describe the manner in which the parameters of the Base Case are varied, in each of the six scenario sets.

The MMFE Model Hurley et al. (2006) described the MMFE model of forecast evolution. In the multiplicative version of the MMFE, for every pair of times $s,t,\ 0 \le s \le t \le T,\ 1 \le t$, there is a forecast D_{st} of the demand that will occur in period t, which was generated at the end of period s, i.e., at the beginning of period $s+1$. The actual demand is $D_t = D_{tt}$, observed at the end of period t. We assume that forecasts are unbiased, so that $D_{st} = E[D_t|\mathcal{F}_{s+1}]$. There is a forecast horizon $H \le T$. The corporate forecasting process generates forecasts that extend H time periods into the future. Therefore, $D_{st} \ne D_{s-1,t}$ if $t < s+H$, because in that case D_{st} was effected by the forecasting process that occurred at the end of period s. However, if $t \ge s+H$ then the end-of-period-s forecasting process did not consider the period-t demand, and $D_{s,t} = D_{s-1,t}$. At the beginning of the time horizon we are given the initial set of forecasts, $\boldsymbol{d_0} = (d_{0,t} : 1 \le t \le T)$. (In this case we use lower case because these forecasts have already been observed). Seasonality and trend are introduced into the model my choosing the vector $\boldsymbol{d_0}$ appropriately.

We model the process by which forecasts are created as follows. The period-s update vector is $\gamma_s = (\gamma_{st} : s \le t < s+H)$. At the end of time period s the update vector $\boldsymbol{\gamma_s}$ is observed, and the multiplicative MMFE model updates forecasts using the formula $d_{st} = \gamma_{st}\, d_{s-1,t}$ for $t = s, s+1, ..., s+H-1$, and by $d_{st} = d_{s-1,t}$ for $t \ge s+H$. In our experiments $\boldsymbol{\gamma_s} = e^{\boldsymbol{\epsilon_s}}$, where the H - dimensional random vector ϵ_s is normally distributed with mean $-\frac{1}{2}\text{diag}(\Sigma_s)$ and variance-covariance matrix Σ_s, and γ_s has a multivariate lognormal distribution whose mean is a vector of ones. $\Sigma_s \sim \Sigma$ and $\gamma_s \sim \gamma$ are both stationary over time.

In the multiplicative MMFE model, it is not hard to show that at the end of period s, given the current information set f_{s+1} and forecast vector $\boldsymbol{d_s} = (d_{s,t} : s \le t \le T)$, the future demands $(D_{tt} : s < t \le T)$ have a conditional distribution that is multivariate lognormal, with easily-computable parameters. Three of our six scenario sets study second order effects, which we create by using different variance-covariance matrices Σ.

The Base Case In the Base Case our holding and backorder costs per unit per period are stationary, equal to $h_t = 1$ and $p_t = 10$. All experiments are conducted for two different lead times: $L = 0$ and $L = 4$. Therefore, to facilitate comparisons between different scenarios, costs are not counted during the first four time periods. Note that when $L = 4$, in the first four time periods the costs incurred are determined by decisions made in the past, and are not influenced by our choice of policy. There is neither trend nor seasonality in the Base Case, so the initial demand forecast is flat, with $\boldsymbol{d_0} = (400, 400, \ldots, 400)$. The time horizon has length $T = 40$, and the horizon over which the user actively generates forecasts has length $H = 12$. This implies that at all times s, the first 13 elements of the forecast vector $\boldsymbol{d_s}$ will be different from each other, but the 13-th element and every subsequent element will be equal to 400.

In the Base Case, we have *constant learning*, meaning that all of the entries on the diagonal of Σ are equal. The diagonal elements are selected so that for $t \ge 12$, the coefficient of variation of the period-t demand D_{tt}, seen from the beginning of time period 1, is 0.75.

The off-diagonal entries of the covariance matrix Σ determine the degree of correlation between the updates that are observed in a given time period, say, time period s. The Base Case assumes that there is some correlation between these updates, modeled by having non-zero, positive values in the first off-diagonal of Σ. Consequently, in the Base Case, if the forecast for the demand in month t will go up in period s (i.e., if $D_{st} > D_{s-1,t}$), then the forecast for demand in month $t+1$ is likely to increase in period s as well (i.e., $P(D_{s,t+1} > D_{s-1,t+1}) > 0.5$). However, $D_{st} > D_{s-1,t}$ does not tell us anything about the forecast $D_{s,t+2}$ for demand in month $t+2$. The values of the non-zero off-diagonal elements are chosen to give a correlation coefficient of 0.5 for each pair of adjacent forecast updates. That is, for each s and each t, $s \le t \le s+H-2$, the update factors γ_{st} and $\gamma_{s,t+1}$ observed in period s have correlation coefficient 0.5, but γ_{st} and $\gamma_{s,t+2}$ are stochastically independent.

Product Launch Scenarios In this set of scenarios we study the effect of rising demand, as might be encountered at a product launch. Again, only the initial forecast vector $\boldsymbol{d_0}$ is varied. For comparison with the base case, we ensure that the mean of the values in $\boldsymbol{d_0}$ is 400. We consider upward demand trends of +5, +10 and +20 per period. In addition, we consider two examples in which the demand rises in a steeper, non-linear manner, mid-way through the horizon; these are generated using an appropriately scaled normal CDF curve. The five initial forecast vectors are plotted in Figure 8.

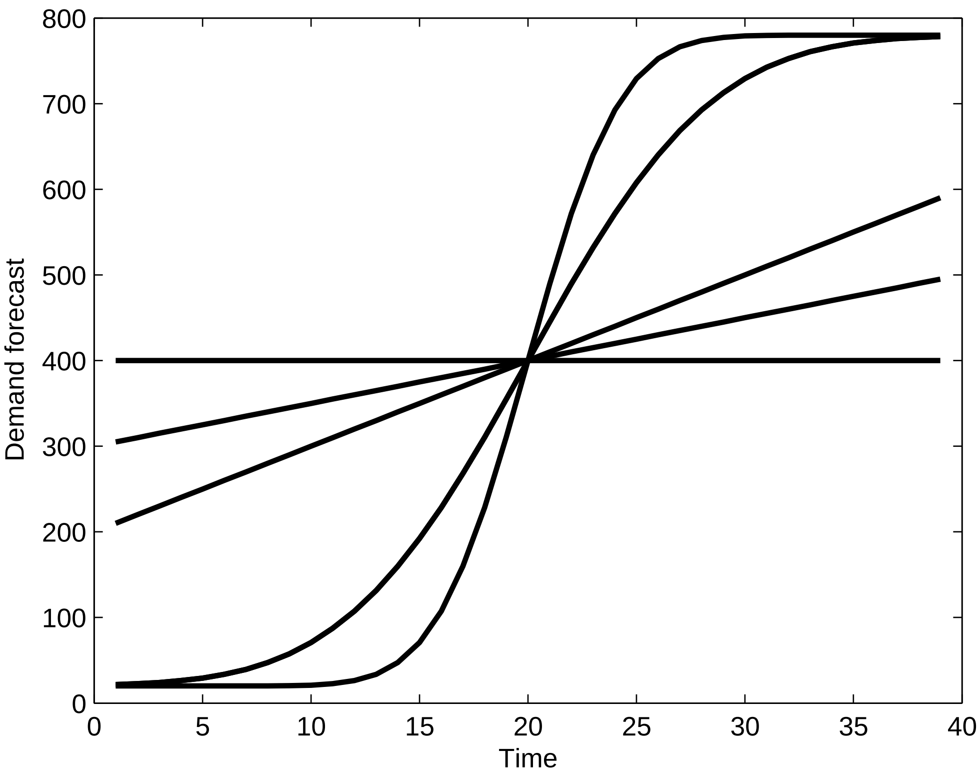

Figure 8 Initial forecast vectors used in Product Launch Scenarios.

End-of-Life Scenarios Here, we study scenarios associated with products that are in an end of life situation, namely those with decreasing initial forecast vectors. Essentially, these are the reverse of the Product Launch scenarios; we have initial forecast vectors with forecasted demand decreasing by $5, 10$ and 20 per period. We also consider two products whose demands have steeper drop-off curves, generated using the normal complementary CDF curve. In addition, we study a total demand crash, in which the demand is forecast to crash to 0 midway through the time horizon.

Seasonality Scenarios In the seasonality study, we use the common base-values described above for all parameters except for the initial forecast vector $\boldsymbol{d_0}$. We conduct experiments with two forms of seasonality, one defined via a sinusoidal function and the other via a step function. In both cases, the maximum value attained is 700 and the minimum is 100. This allows us to compare results more easily with the base case, because the mean of the entries in the initial forecast vector is 400 in all cases.

By the *cycle length*, we mean the number of time periods between two consecutive high-points. We consider cycle lengths with values 2, 4 and 8. For example, for the step-function with period 4, we have $\boldsymbol{d_0} = (700, 700, 100, 100, 700, 700, 100, 100, \dots)$.

The above scenario sets test the effect of varying $\boldsymbol{d_0}$, the initial forecast vector. In the final three scenario sets, we focus instead on varying Σ. In all of these, we take $\boldsymbol{d_0} = (400, 400, \dots, 400)$.

Coefficient of Variation Scenarios In this scenario set, we study the effect of varying the magnitude of the variance in the demands and the forecasts. Note that for $t \geq H = 12$, at the end of time period $t - H$, we have $D_{tt} = \Gamma_t \ d_{t-H,t}$, where Γ_t is random and has the same distribution as

$$\Gamma = \Gamma_H = \Pi_{i=1}^{H} \gamma_{H+1-i,H} = \exp\left(\sum_{i=1}^{H} \epsilon_{H+1-i,H}\right). \tag{4}$$

The $\epsilon_{H+1-i,H}$'s are independent normal random variables, with mean such that $E[e^{\epsilon_{H+1-i,H}}] = 1$, and with variance σ_{ii}, the i-th diagonal element of Σ, our forecast update matrix. (Note that σ_{ii} is a variance, not a standard deviation). Thus, Γ is log-normal, with mean one and variance $e^{(\Sigma_{i=1}^{H}\sigma_{ii})} - 1$. The coefficient of variation of Γ, and of D_{tt} for $t \geq H$, is $[e^{(\Sigma_{i=1}^{H}\sigma_{ii})} - 1]^{1/2}$. In the Base Case this number is 0.75. In the scenarios where we investigate the effect of variance, we scale the entries of Σ such that the coefficient of variation of Γ takes specific values, namely 0.5, 0.7, 1, 2, 4, and 8. This corresponds to different levels of demand variability.

Time of Learning Scenarios If $s \leq t \leq s + H$ then the logic behind equation (4) above indicates that at the end of period s, the random variable $(D_{tt}|d_{st})$ has mean 1 and variance $e^{(\Sigma_{i=1}^{t-s}\sigma_{ii})} - 1$. Therefore we use $(\sum_{i=1}^{t-s}\sigma_{ii})$, which in the Base Case ranges from 0 to 0.446, to measure the portion of the total variability in D_{tt} that is still *unresolved* in period s. In Figure 9 we plot $(\sum_{i=1}^{t-s}\sigma_{ii})$ as a function of $t - s$, for $0 \leq t - s \leq H$. The different curves represent four different possibilities for the way in which variability is resolved. In the Base Case we have *constant learning*, meaning that all of the entries in diag(Σ) are equal, and the curve is a straight line. When the diagonal of Σ has relatively large values in the lower right portion of the matrix, the plot is convex, and the unresolved uncertainty is low when s is close to t. This corresponds to *early learning*. Conversely, when the values in the diagonal of Σ are weighted towards the upper right corner of the matrix we have *late learning*, the plot is concave, and most of the uncertainty about the true value of D_{tt} is resolved in periods s that are close to t. We also consider the setting in which there is more weight in the center of the diagonal of Σ than at the extremes. In this case most of the learning takes place near the middle of the forecast horizon.

We construct variance-covariance matrices Σ to correspond with these four cases: constant, early, late and mid-horizon learning. In all cases, the values of Σ are scaled to ensure that the coefficient of variation of Γ, and of D_{tt} for $t \geq 12$, remain constant at 0.75.

Correlation Scenarios In this scenario set we test the effect of different types of correlation between the updates. We vary correlation in two ways. First, we set the number of non-zero off-diagonals of our 12x12 matrix, Σ, to 0 (which corresponds to no correlation), 1, 4 and 8. Secondly, the sign of the off-diagonal elements can be all positive, all negative, or entries alternating between positive and negative. (The base case corresponds to 1 off-diagonal with non-zero elements which are all positive.) As in the base case, the diagonal elements of Σ are all equal (the constant learning case), and the coefficient of variation of Γ is 0.75.

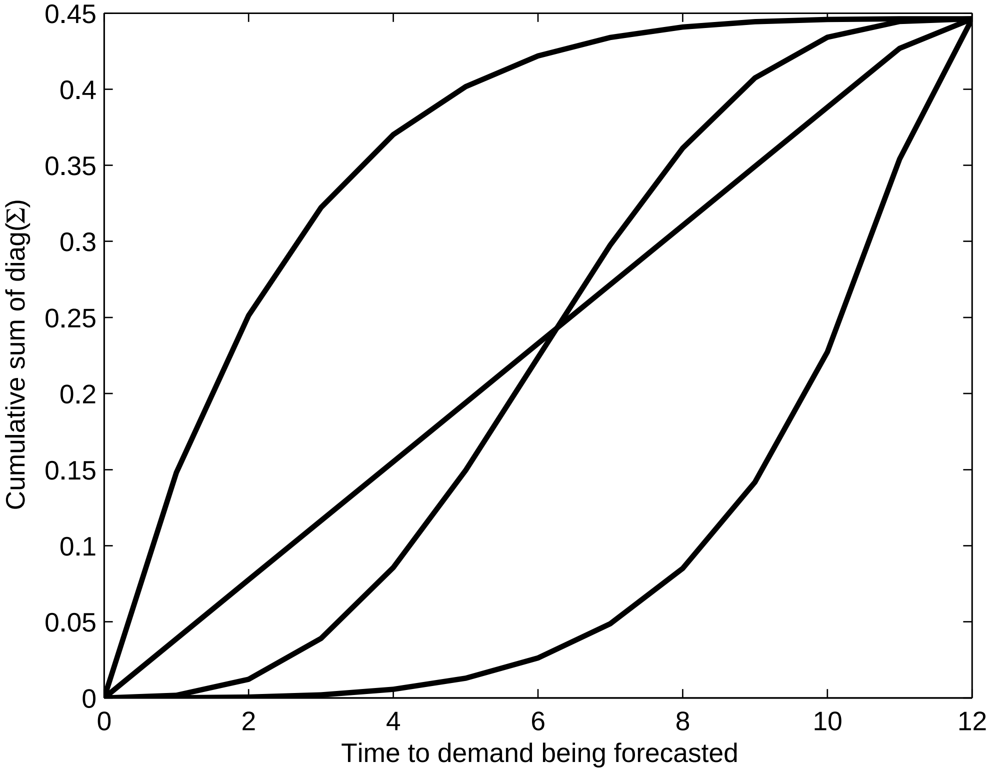

Figure 9 Cumulative sum of the diagonal elements of Σ for constant, early, late and mid-horizon learning.

Table 1 summarizes the scenarios we study. The number of scenarios for each set is given in parentheses after the set name; we see that there are 38 in total. We run each of these with lead times $L = 0, 4$. In addition, the 6 seasonality-based scenarios were run with $L = 8$. That makes a total of $38 \times 2 + 6 = 82$ scenario - lead time pairs. For each of the pairs, we ran $N = 1,000$ independent trials.

Appendix C: Extensions of the Dual-Balancing Policy

Integer-Valued Demands We now discuss the case in which the demands are integer-valued random variables, and the order in each period is also restricted to an integer. A simple, illustrative example of the Dual-Balancing policy with integer-valued demands is found in Appendix A. In the integer-valued demand case, in each period s, the functions $H_s^B(q_s^B)$ and $\tilde{\Pi}_s^B(q_s^B)$ are originally defined only for integer values of q_s^B. We now define these functions for any value of q_s^B by interpolating piecewise linear extensions of the integer values. It is clear that these extended functions preserve the convexity and monotonicity properties discussed in the previous (continuous) case. However, it is still possible (and even likely) that the value q_s' that balances the functions l_s^B and $\tilde{\pi}_s^B$ is not an integer. Instead we consider the two consecutive integers q_s^1 and $q_s^2 := q_s^1 + 1$ such that $q_s^1 < q_s' < q_s^2$. In particular, $q_s' := \lambda q_s^1 + (1 - \lambda) q_s^2$ for some $0 < \lambda < 1$. We now order q_s^1 units with probability λ and q_s^2 units with probability $1 - \lambda$. This constructs what we call a *randomized Dual-Balancing policy.*

Observe that at the beginning of time period s the order quantity of the Dual-Balancing policy is still a random variable $Q_s^B = Q_s'$ with support $\{q_s^1, q_s^2\} = \{q_s^1(f_s), q_s^2(f_s)\}$, which is a function of the

Table 1 Scenario codes

Topic	Description
Product Launch (5)	Increment by I per period, $I \in \{5, 10, 20\}$ Increasing scaled normal CDF curve Steeper increasing scaled normal CDF curve
End-of-Life (6)	Decrement by I per period, $I \in \{5, 10, 20\}$ Decreasing scaled normal CDF curve Steeper decreasing scaled normal CDF curve Demand crash
Seasonal (7)	Initial forecast vector is flat Sinusoidal periodicity with cycle length n, $n \in \{2, 4, 8\}$ Step-function periodicity with cycle length n
Coeff. of Var. (6)	Coefficient of variation equals $\beta = 0.5, 0.7, 1, 2, 4, 8$
Learning Rate (4)	Constant learning Late learning Early learning Mid-horizon learning
Correlation (10)	All off-diagonal elements of Σ are 0 First n off-diagonals of Σ have positive entries, $n \in \{1, 4, 8\}$ First n off-diagonals of Σ have negative entries First n off-diagonals of Σ have entries alternating positive and negative

observed information set f_s. We would like to show that this policy admits the same performance guarantee of 2. For each $t = 1, \ldots, T - L$, let Z_t be again the random balanced cost of the Dual-Balancing policy in period t. Focus now on some period s. For a given observed information set $f_s \in \mathcal{F}_s$ we have for some $0 \le \lambda = \lambda(f_s) \le 1$,

$$z_s = E[H_s^B(Q'_s)|f_s] = \lambda E[H_s^B(q_s^1)|f_s] + (1-\lambda)E[H_s^B(q_s^2)|f_s] := E[H_s^B(\lambda q_s^1 + (1-\lambda)q_s^2)|f_s],$$

and

$$z_s = E[\tilde{\Pi}_s^B(Q'_s)|f_s] = \lambda E[\tilde{\Pi}_s^B(q_s^1)|f_s] + (1-\lambda)E[\tilde{\Pi}_s^B(q_s^2)|f_s] := E[\tilde{\Pi}_s^B(\lambda q_s^1 + (1-\lambda)q_s^2)|f_s].$$

The second equality (in each of the two expressions above) is a formal statement of the fact that we extended the domains of $H_s^B(q_s^B)$ and $\tilde{\Pi}_s^B(q_s^B)$ from integer to real values using piecewise linear interpolation. By the definition of the algorithm we have,

$$\lambda E[H_s^B(q_s^1)|f_s] + (1-\lambda)E[H_s^B(q_s^2)|f_s] = \lambda E[\tilde{\Pi}_s^B(q_s^1)|f_s] + (1-\lambda)E[\tilde{\Pi}_s^B(q_s^2)|f_s].$$

It is now readily seen that, for each period s and each $f_s \in \mathcal{F}_s$, we again have $E[H_s^B(Q'_s) + \tilde{\Pi}_s^B(Q'_s)|f_s] = 2z_s$, i.e., $E[H_s^B(Q'_s) + \tilde{\Pi}_s^B(Q'_s)|\mathcal{F}_s] = 2Z_s$. This also implies that Lemma 1 is still valid.

Now define the sets $\mathcal{T}_H$ and $\mathcal{T}_\Pi$ in the following way. Let $\mathcal{T}_H = \{t : X_t^B + Q_t^2 \le Y_t^{OPT}\}$, and $\mathcal{T}_\Pi = \{t : X_t^B + Q_t^2 > Y_t^{OPT}\}$. Observe that for each period s, conditioned on some $f_s \in \mathcal{F}_s$, we know deterministically x_s^B, q_2^B and, if the optimal policy is deterministic, we also know y_s^{OPT}. Therefore, we know whether $s \in T_H$ or $s \in \mathcal{T}_\Pi$. If the optimal policy is also a randomized policy, we condition not only on f_s but also on the decision made by the optimal policy in period s. Moreover, if $s \in T_H$, then, with probability 1, $Y_s^B \le Y_s^{OPT}$, and if $s \in \mathcal{T}_\Pi$, then, with probability 1, $Y_s^B \ge Y_s^{OPT}$. This implies that also Lemmas 2 and 3 are still valid. The following theorem is now established (the proof is identical to that of Theorem 2 above).

Theorem C.1 The randomized Dual-Balancing policy has a worst-case performance guarantee of 2, i.e., for each instance of the capacitated periodic-review stochastic inventory control problem, the expected cost of the randomized Dual-Balancing policy is at most twice the expected cost of an optimal solution, i.e, $E[\mathcal{C}(B)] \leq 2E[\mathcal{C}(OPT)]$.

Stochastic Lead Times In this section, we consider the more general model where the lead time of an order placed in period s is some nonnegative integer-valued random variable L_s. However, we assume that the random variables $L_1, \dots, L_T$ are correlated, and in particular, that $s + L_s \leq t + L_t$ for each $s \leq t$. In other words, we assume that any order placed at time s will arrive no later than any other order placed after period s. This is a very common assumption in the inventory literature, usually described as "no order crossing".

Similar to Levi et al. (2007), we next describe how to extend the Dual-Balancing policy and the analysis of the worst-case expected performance to this more general setting. For each $t = 1, \dots, T$, let S_t be the latest period for which an order placed in that period arrives before time t. In other words, $S_t := \max\{s : s + L_s \leq t\}$. Now modify the definition of the random variables W_{st} (for each $s \leq t$) to be

$$W_{st} := \min\{\mathbb{1}(s \leq S_t)\bar{Q}_s, \mathbb{1}(s \leq S_t)\ (D(_{[s,t]} - (X_s + Q_s + \sum_{j \in (s, S_t]} u_j))^+\}.$$

Similar to the discussion in Section 3 above, we can write

$$W_{st} = \mathbb{1}(s \leq S_t) \left[(D_{[s,t]} - (X_s + Q_s + \sum_{j \in (s, S_t]} u_j))^+ - (D_{[s,t]} - (X_s + \sum_{j \in [s, S_t]} u_j))^+ \right],$$

and

$$W_{st} = \mathbb{1}(s \leq S_t) \left[(D_t - NI_t - \sum_{j \in (s, S_t]} \bar{Q}_j)^+ - (D_t - NI_t - \sum_{j \in [s, S_t]} \bar{Q}_j)^+ . \right]$$

We again define the forced marginal backlogging cost in period s as $\tilde{\Pi}_s = \sum_{t \geq s} p_t W_{st}$. It is straightforward to check that we can still express the cost of each feasible policy P as $\mathcal{C}(P) = \sum_t (H_t + \tilde{\Pi}_t)$. In each period, we again balance the conditional expected marginal holding cost against the conditional expected forced marginal backlogging cost. It is readily verified that the same analysis described in Section 4.1 is still valid.

Theorem C.2 The Dual-Balancing policy provides a performance guarantee of 2 for the capacitated periodic-review stochastic inventory control problem with stochastic lead times and non-crossing orders.

Acknowledgments

We thank the Associate Editor and the anonymous referees for numerous comments that improve the content and exposition of this paper.

This research was partially conducted while the first and the fourth authors were a PhD students in the ORIE department at Cornell University.

The research of the first author was supported partially by NSF grants CCR-9912422 & CCR-0430682.

The research of the second author was supported partially by NSF grants DMI-0075627 & DMI-0500263, and the Querétaro Campus of the Instituto Tecnológico y de Estudios Superiores de Monterrey. The research of the third author was supported partially by NSF grants CCR-9912422, CCR-0430682, & DMI-0500263.

The research of the fourth author was supported partially by NSF grant DMI-0500263

References

Aviv, Y., A. Federgruen. 1997. Stochastic inventory models with limited production capacity and periodically varying parameters. *Probability in the Engineering and Informational Sciences* **11** 107–135.

Axsäter, S. 1990. Simple solution procedures for a class of two-echelon inventory problems. *Operations Research* **38** 64–69.

Chan, E. W. M. 1999. Markov chain models for multi-echelon supply chains. Ph.D. thesis, School of OR&IE, Cornell University, Ithaca, NY.

Chen, F., J.S. Song. 2001. Optimal policies for multi-echelon inventory problems with Markov-modulated demand. *Operations Research* **49(2)** 226–234.

Dong, L., H. L. Lee. 2003. Optimal policies and approximations for a serial multiechelon inventory system with time-correlated demand. *Operations Research* **51(6)** 969–980.

Federgruen, A., P. Zipkin. 1986a. An inventory model with limited production capacity and uncertain demands i: The average-cost criterion. *Mathematics of Operations Research* **11(2)** 193–207.

Federgruen, A., P. Zipkin. 1986b. An inventory model with limited production capacity and uncertain demands ii: The discounted-cost criterion. *Mathematics of Operations Research* **11(2)** 208–215.

Glasserman, P. 1997. Bounds and asymptotics for planning critical safety stocks. *Operations Research* **45** 224–257.

Heath, D. C., P. L. Jackson. 1994. Modeling the evolution of demand forecasts with application to safety stock analysis in production distribution-systems. *IIE Transactions* **26**(3) 17–30.

Hurley, Gavin, P. Jackson, R. Levi, R. O. Roundy, D. B. Shmoys. 2006. New policies for stochatic inventory control models - a theoretical and computational study. Working paper.

Ignall, E., A. F. Veinott. 1969. Optimality of myopic inventory policies for several substitute products. *Managment Science* **15** 284–304.

Iida, T., P. Zipkin. 2001. Approximate solutions of a dynamic forecast-inventory model. Working paper.

Janakiraman, G., J. A Muckstadt. 2003. A decomposition approach for a capacitated singel-stage inventory system. Tech. Rep. TR1360, ORIE Department, Cornell University.

Kapuscinski, R., S. Tayur. 1998. A capacitated production-inventory model with periodic demand. *Operations Research* **46(6)** 899–911.

Katircioglu, K., D. Atkins. 1996. New optimal policies for a unit demand inventory problem. Unpublished manuscript.

Levi, R., Martin Pál, R. O. Roundy, D. B. Shmoys. 2007. Approximation algorithms for stochastic inventory control models. *Mathematics of Operations Research* **32** 284–302.

Lu, X., J. S. Song, A. C. Regan. 2006. Inventory planning with forecast updates: approximate solutions and cost error bounds. *Operations Research* **54** 1079–1097.

Metters, R. 1997. Production planning with stochastic seasonal demand and capacitated production. *IIE Transactions* **29** 1017–1029.

Muharremoglu, A., J. N. Tsitsiklis. 2001. A single-unit decomposition approach to multi-echelon inventory systems. Working paper.

Roundy, R. O., J. A. Muckstadt. 2000. Heuristic computation of periodic-review base stock inventory policies. *Management Science* **46(1)** 104–109.

Song, J.S, P. Zipkin. 1993. Inventory control in a fluctuating demand environment. *Operations Research* **41(2)** 351–370.

Tayur, S. 1992. Computing the optimal policy for capacitated inventory models. *Stochastic Models* **9** 585–598.

Veinott, A. F. 1965. Optimal policy for a multi-product, dynamic, non-stationary inventory problem. *Management Science* **12** 206–222.

Zipkin, P. H. 2000. *Foundations of inventory management*. The McGraw-Hill Companies, Inc.