INFORMS TutORials in Operations Research

Publication details, including instructions for authors and subscription information:
http://pubsonline.informs.org

Provably Near-Optimal Approximation Algorithms for Operations Management Models

Retsef Levi

To cite this entry: Retsef Levi. Provably Near-Optimal Approximation Algorithms for Operations Management Models. *In* INFORMS TutORials in Operations Research. Published online: 14 Oct 2014; 179-192.
https://doi.org/10.1287/educ.1100.0076

TUTORIALS IN
OPERATIONS RESEARCH
INFORMS 2010

 | ISBN 978-0-9843378-0-4
DOI 10.1287/educ.1100.0076

Provably Near-Optimal Approximation Algorithms for Operations Management Models

Retsef Levi
Sloan School of Management, Massachusetts Institute of Technology,
Cambridge, Massachusetts 02139, retsef@mit.edu

Abstract Multistage stochastic optimization models arise in various application domains, including central areas of operations management, such as inventory control, supply chain management, and revenue management. Unfortunately, it is usually computationally very hard to find optimal solutions for these fundamental models, and typically even finding good solutions is very challenging, both in theory and practice. In this tutorial, we will survey various relatively new algorithmic and performance analysis techniques that can be applied to some of these models to construct provably near-optimal algorithms, particularly algorithms that admit worst-case performance guarantees. These techniques span ideas from various disciplines, such as optimization, computer science, and stochastic analysis.

Keywords approximation algorithms; inventory control; revenue management; stochastic control

1. Introduction

Many if not most of the core problems studied in operations management fall into the category of *multistage stochastic optimization models*. Particularly, one has to make multiple, typically dependent decisions over time to optimize a certain objective function under uncertainty on how the system will evolve over the future time horizon. Unfortunately, it is usually computationally very hard to find optimal solutions for these fundamental models, and typically even finding good solutions is very challenging, both in theory and practice. Thus, in many scenarios there is a challenging gap of leveraging the respective models being used into effective and practical policies.

In this tutorial, we will survey several recent algorithmic and performance analysis techniques that can be applied to some of these models to construct provably near-optimal algorithms, particularly algorithms that admit *worst-case performance guarantees*. The notion of worst-case performance guarantees has been used extensively in computer science in the analysis of *approximation algorithms* for combinatorial NP-hard problems (Vazirani [39]). Put formally, an algorithm is called an α-approximation algorithm or is said to have a *worst-case guarantee* of α (for some constant $\alpha > 1$) if it is a polynomial time algorithm, and for any instance of the problem the algorithm is guaranteed to provide a solution with cost that is at most α times the optimal cost. (The definition applies to minimization problems; there is a straightforward equivalent definition for maximization problems.) Intuitively, the notion of a worst-case guarantee could be viewed in the context of a "game" between the algorithm designer and an adversary opponent. The algorithm designer proposes an algorithm, and the adversary attempts to generate the worst instance, for which the relative gap of the cost of the proposed algorithm and the optimal cost is maximized. The worst-case guarantee is simply a statement to what extent the adversary can be successful with respect to a given algorithm.

Traditionally, approximation algorithm techniques have been applied primarily to deterministic combinatorial optimization problems. The work on approximation algorithms for stochastic combinatorial problems goes back to the work on stochastic scheduling problems of Möhring et al. [29, 30] and the more recent work of Möhring et al. [31]. Recently, there has been a growing stream of approximation results for several two-stage stochastic combinatorial problems. For a comprehensive literature review, we refer the reader to Stougie and van der Vlerk [37], Dye et al. [9], and Shmoys and Swamy [34, 35]. In contrast, this tutorial is focused on the relatively harder multistage stochastic optimization models, for which there has been relatively little work (for example, see Dean et al. [8], Swamy and Shmoys [38], Chan and Farias [5]).

The concept of approximation algorithms has been applied to several problems in operations management, but again primarily to deterministic problems; for examples, see Silver and Meal [36], Roundy [32], and Levi et al. [24, 26, 22]. Until recently, there have been relatively few examples of worst-case analysis of heuristics for stochastic optimization models within operations management (Chen [6]). In fact, with relatively few exceptions (e.g., Gallego and van Ryzin [10], Lu et al. [28], Halman et al. [11], Chu and Shen [7]), most of the heuristics and algorithms that have been proposed for operations management models were tested merely through computational experiments on randomly generated instances. This does not necessarily provide strong indications that the proposed heuristics are good in general, beyond the instances that were actually tested. In contrast, worst-case performance analysis has the advantage that it provides a priori and posteriori guarantees on the quality of the solution produced by the algorithm. Moreover, the performance analysis provides insights on how to design algorithms that have good typical (empirical) performance, which in most cases is significantly better than the worst-case analysis. We note that there has been recent work on near-optimal and asymptotically optimal heuristics for stochastic optimization models with learning, but this is not within the scope of this survey.

In this tutorial, we will discuss some of the recent work to develop provably near-optimal approximation algorithms for operations management models. We shall describe the respective algorithms and their theoretical (worst-case) and typical (computational) performance analysis. In addition, we shall highlight some of the central techniques that have been used, and point out interesting future research directions. As will be demonstrated, the respective techniques span ideas from many disciplines, such as optimization, computer science, and stochastic analysis. The discussion in this tutorial is focused on two classes of models, specifically, core stochastic inventory control models (following primarily Levi et al. [27, 25]) and revenue management models of reusable resources (following closely Levi and Radovanovic [19]). Although the focus of the tutorial is on operations management models, the techniques that will be discussed have a promising potential to be applicable in other application domains.

2. Stochastic Inventory Control Models

This section follows closely the discussion in Levi et al. [25, 27] on core stochastic inventory control models. These models have been studied extensively in the inventory management literature, primarily based on the paradigm of *dynamic programming*. Dynamic programming approaches have been affective in characterizing the structure of the optimal policies for some of these models. However, in most practical cases, the resulting dynamic programs are computationally very challenging because of the fact that the respective state space grows large. Thus, they usually do not lead to effective algorithmic approaches.

2.1. Model Formulation

We consider a finite planning horizon of T periods numbered $t = 1, \ldots, T$ with stochastic demands, denoted by $D_1, \ldots, D_T$, respectively. At the beginning of each period s, there is an observed *information set* that is denoted by f_s. The information set contains all of the information that is available by the beginning of time period s. More specifically, the information set f_s consists of the realized demands $(d_1,\ldots,d_{s-1})$ over the interval $[1,s)$, and possibly some more (external) information denoted by $(w_1,\ldots,w_s)$. The information set f_s in period s is one specific realization in the set of all possible realizations of the random vector $F_s=(D_1,\ldots,D_{s-1},W_1,\ldots,W_s)$. This set is denoted by $\mathcal{F}_s$. The joint conditional distribution of future demands $(D_s,\ldots,D_T)$ as seen from period s is determined by the observed information set f_s. The only assumption on this distribution is that for each $f_s\in\mathcal{F}_s$, the conditional expectation $\mathrm{E}[D_t \mid f_s]$ is well defined and finite for each period $t\geq s$. In particular, this captures all of the known forecasting models and allows correlation between demands in different periods.

At the beginning of each period s, an order of size $q_s\geq 0$ is placed; this order will arrive and become available only after a lead time of L periods. Only then is the actual demand in period s observed and satisfied to the maximum extent possible. In each period $t=1,\ldots,T$, several 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 *fixed ordering cost*, K_t, that is incurred whenever a positive order is placed regardless of its size, a per-unit *holding cost*, $h_t\geq 0$, for holding excess inventory from period t to $t+1$, and a per-unit *backlogging penalty cost*, $p_t\geq 0$, that is incurred for each unsatisfied unit of demand at the end of period t. Unsatisfied units of demand are assumed to be backlogged; that is, each unit of unsatisfied demand stays in the system until it is satisfied. The goal is to satisfy all demands with minimum total expected cost. (There is also a discount factor β that, without loss of generality, is assumed to be 1.)

Next we shall discuss the *uncapacitated periodic-review stochastic inventory control problem*, which is the special case of the above model, assuming there are no fixed costs (i.e., $K_t=0$, for each t) and no capacity constraints (i.e., $u_t=\infty$, for each t). In addition, assume (almost) without loss of generality (see Levi et al. [25]) that the per-unit ordering costs are all equal to 0 (i.e., $c_t=0$, for each period t). As already mentioned, the traditional approach to study these models has been dynamic programming. Using the dynamic programming approach it can be shown that *state-dependent base-stock policies are optimal* (Zipkin [42]). However, the computational complexity of the resulting dynamic programs is very sensitive to the dimension of the sets $\mathcal{F}_s$. In particular, in many practical scenarios these sets are of high dimension, which leads to dynamic programming formulations that are computationally intractable. (For an elaborated discussion, see Levi et al. [25].) In fact, it has been shown in Halman et al. [11] that this model is $\#P$-hard, even for the relatively simple special case of independent discrete, finite support demands. In fact, for this special case it is possible to construct a polynomial approximation scheme (PTAS); that is, the problem can be approximated to an arbitrary degree of accuracy with a running time that depends on the degree of accuracy.

2.2. Dual-Balancing Policies

In this section, we shall describe the *dual-balancing* policy following Levi et al. [25]. This policy is based on two major ideas:

Marginal cost accounting scheme. The standard dynamic programming approach directly assigns to the decision of how many units to order in each period only the expected holding and backlogging costs incurred in that period, although this decision might effect the costs in future periods. Instead, marginal cost accounting scheme assigns to the decision in each period *all* the expected costs that, once this decision is made, become unaffected by any decision made in future periods. (These cost may still depend on future demands.)

Cost balancing. The idea of cost balancing was used in the past to construct heuristics with constant performance guarantees for deterministic inventory problems (see, for example, Silver and Meal [36]). The key observation in the above model in is that any policy in any period incurs potential expected costs due to *over ordering* (namely, expected holding costs of carrying excess inventory) and *under ordering* (namely, expected backlogging costs incurred when demand is not met on time). The dual-balancing policy described below simply aims to repeatedly balance the expected (marginal) holding cost against the expected (marginal) backlogging cost. The cost balancing depends on the marginal cost accounting scheme.

2.2.1. Marginal Cost Accounting Scheme. We first introduce a *marginal holding cost accounting approach.* Without loss of generality, assume that the ordered supply units are consumed on *first-ordered, first-consumed basis.* The key observation under this assumption is that once an order is placed in some period, then the expected holding cost that the units just ordered will incur over the rest of the planning horizon is a function only of the realized demands over the rest of the horizon, not of any future orders. Hence, with each period, we can associate the overall expected holding cost that is incurred by the units ordered in this period over the entire horizon. We note that similar ideas of holding cost accounting were used previously in the context of models with continuous time, infinite horizon, and stationary (Poisson distributed) demand (see, for example, the work of Axsäter and Lundell [3], Axsäter [2]). More specifically, let x_s be the *inventory position* at the beginning of period s that captures the total sum of the physical *on-hand inventory* and the *outstanding orders* (placed in past periods, but still on the way) minus the pending backlogged demand. Say now that q_s units were ordered in period s, and consider a future period $t \geq s$. Then the holding cost incurred by the q_s at the end of period t is $h_t(q_s - (D_{[s,\,t]} - x_s)^+)^+$, where $x^+ = \max(x, 0)$ and $D_{[s,\,t]} = \sum_{j=s}^{t} D_j$ is the cumulative demand over the interval $[s,t]$. Observe that if $D_{[s,\,t]} \leq x_s$, then none of the q_s units has been yet consumed. When $D_{[s,\,t]}$ exceeds x_s, the q_s units are used to satisfy the demand until all of them are consumed. It follows that the total holding cost incurred by the q_s units ordered in period s over the entire horizon is equal to

$$H_s = H_s(Q_s) := \sum_{t=s+L}^{T} h_t\big(Q_s - (D_{[t,\,s]} - X_s)^+\big)^+. \tag{1}$$

Because X_s and Q_s are realized at the beginning of period s (whereas, x_s and q_s are the realizations of X_s and Q_s, respectively), then, as seen from the beginning of period s, this quantity depends only on future demands and not on any future decision.

In addition, in an uncapacitated model the decision of how many units to order in each period affects the expected backlogging cost in only a single future period, namely, a lead time ahead. Now let Π_s be backlogging cost incurred in period $s+L$, for each $s = 1-L, \ldots, T-L$. In particular, it is straightforward to verify that

$$\Pi_s := p_{s+L}\big(D_{[s,\,s+L]} - (X_s + Q_s)\big)^+, \tag{2}$$

where $D_j := 0$ with probability 1 for each $j \leq 0$. (Observe that the supply units captured by $X_s + Q_s$ will become available by time period $s+L$, and that no order placed after period s will arrive by time period $s+L$.)

Now let $\mathcal{C}(P)$ be the cost of a feasible policy P and use the superscript P to relate the respective quantities to that policy. Clearly,

$$\mathcal{C}(P) := \sum_{t=1-L}^{0} \Pi_t^P + H_{(-\infty,\,0]} + \sum_{t=1}^{T-L} (H_t^P + \Pi_t^P), \tag{3}$$

where $H_{(-\infty,\,0]}$ denotes the total holding cost incurred by units ordered before period 1 (given as an input). We note that the first two expressions $\sum_{t=1-L}^{0} \Pi_t^P$ and $H_{(-\infty,\,0]}$ are the same for any feasible policy and each realization of the demand, and therefore we will omit them. Because they are nonnegative, this will not affect our approximation results. Also observe that, without loss of generality, it can be assumed that $Q_t^P = H_t^P = 0$ for any policy P and each period $t = T - L + 1, \ldots, T$, because nothing that is ordered in these periods can be used within the given planning horizon. We now can write

$$\mathcal{C}(P) = \sum_{t=1}^{T-L} (H_t^P + \Pi_t^P). \tag{4}$$

The cost accounting scheme in (4) above is marginal; i.e., in each period we account all the expected costs that become unaffected by any future decision.

2.2.2. Dual-Balancing Policy. For simplicity of exposition, the discussion is focused on the case where fractional orders are allowed. In each period $s = 1, \ldots, T - L$, consider the observed information set f_s (where again $f_s \in \mathcal{F}_s$) and x_s^B, the resulting inventory position of the dual-balancing policy at the beginning of period s. The order quantity of the dual-balancing policy is computed based on the two following functions:

(i) the expected holding cost over $[s, T]$ incurred by the *additional* q_s units ordered in period s, conditioned on f_s (we denote this function by $l_s^B(q_s)$, where $l_s^B(q_s) := \mathrm{E}[H_s^B(q_s) \mid f_s]$);

(ii) the expected backlogging cost incurred in period $s+L$ as a function of the additional q_s units ordered in period s, conditioned again on f_s (we denote this function by $\pi_s^B(q_s)$, where $\pi_s^B(q_s) := \mathrm{E}[\Pi_s^B(q_s) \mid f_s]$).

It is easy to verify that these are only functions of q_s. The dual-balancing policy orders $q_s^B = q_s'$ units in period s such that $l_s^B(q_s') = \pi_s^B(q_s')$, i.e., $\mathrm{E}[H_s^B(q_s') \mid f_s] = \mathrm{E}[\Pi_s^B(q_s') \mid f_s]$; that is, q_s' is set to balance the conditional expected marginal over the time interval $[s, T]$ and the conditional expected backlogging cost in period $s + L$. Because fractional orders are allowed, the functions $l_t^B(q_t)$ and $\pi_t^B(q_t)$ are continuous in q_t, for each $t = 1, \ldots, T - L$. Moreover, the function $l_s^B(q_s)$ is equal to 0 for $q_s = 0$ and is an increasing function in q_s, which goes to infinity as q_s goes to infinity. In addition, the function $\pi_s^B(q_s)$ is nonnegative for $q_s = 0$ and is a decreasing function in q_s, which goes to 0 as q_s goes to infinity. Thus, q_s' is well defined. Moreover, q_s' can be computed efficiently using bisection search.

2.2.3. Worst-Case Analysis. Next we shall 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 above, and amortize the period cost of the dual-balancing policy with the cost of the optimal policy; that is, we will show that the optimal policy must incur on expectation at least half of the expected cost of the dual-balancing policy.

Using the marginal holding cost accounting scheme discussed above, the expected cost of the dual-balancing policy can be expressed as

$$\mathrm{E}[\mathcal{C}(B)] = \sum_{t=1}^{T-L} \mathrm{E}[H_t^B + \Pi_t^B]. \tag{5}$$

For each $t = 1, \ldots, T - L$, let Z_t be the *random balanced cost* incurred by the dual-balancing policy in period t, i.e., $Z_t = \mathrm{E}[H_t^B \mid F_t] = \mathrm{E}[\Pi_t^B \mid F_t]$. Note that Z_t is realized in period t as a function of the observed information set f_t (we will denote its realization by z_t). By the construction of the dual-balancing policy, we know that, with probability 1, $\mathrm{E}[H_t^B \mid F_t] = \mathrm{E}[\Pi_t^B \mid F_t]$, for each period $t = 1, \ldots, T - L$. This implies that, for each period t, $\mathrm{E}[H_t^B + \Pi_t^B \mid F_t] = 2Z_t$, which implies the following lemma.

Lemma 1 (Lemma 4.1 in Levi et al. [25]). *The expected cost of the dual-balancing policy is equal to twice the expected sum of the Z_t variables, i.e.,* $\mathrm{E}[\mathcal{C}(B)] = 2\sum_{t=1}^{T-L} \mathrm{E}[Z_t]$.

To complete the analysis it is sufficient to show that the expected cost of an optimal policy is at least $\sum_{t=1}^{T-L} \mathrm{E}[Z_t]$. For each period $t = 1, \ldots, T-L$, let Y_t^P be the inventory position of a policy P in period t *after* ordering, i.e., $Y_t^P = X_t^P + Q_t^P$. Consider the following random partition of the periods $t = 1, \ldots, T-L$ that is induced by comparing the inventory position after ordering of the dual-balancing policy and the optimal policy, denoted by OPT. Let $\mathcal{T}_H$ be the set of periods in which the optimal policy had more inventory than the dual-balancing policy, i.e., the set of periods t, such that $Y_t^B < Y_t^{OPT}$. Let $\mathcal{T}_\Pi$ be the set of periods in which the dual-balancing policy had at least as much inventory as OPT, i.e., the set of periods t such that $Y_t^B \geq Y_t^{OPT}$. In each period $t \in \mathcal{T}_H$, the inequality $Y_t^B < Y_t^{OPT}$ holds and implies that the Q_s^B units ordered by the dual-balancing policy in period t have been ordered by OPT either in period t or even earlier. Thus, the holding cost they incur under OPT are higher than those incurred under the dual-balancing policy. This implies the following lemma.

Lemma 2 (Lemma 4.2 in Levi et al. [25]). *For each realization $f_T \in \mathcal{F}_T$, the marginal holding cost incurred by the dual-balancing policy in all periods $t \in \mathcal{T}_H$ is at most the overall holding cost incurred by OPT, denoted by H^{OPT}, i.e., $\sum_{t \in \mathcal{T}_H} H_t^B \leq H^{OPT}$ with probability* 1.

On the other hand, in each period $t \in \mathcal{T}_\Pi$, the inequality $Y_t^B \geq Y_t^{OPT}$ holds and implies that the backlogging incurred by OPT at the end of period $t+L$ will be higher than that of the dual-balancing policy in that period. The following lemma then follows.

Lemma 3 (Lemma 4.3 in Levi et al. [25]). *For each realization $f_T \in \mathcal{F}_T$, the marginal backlogging penalty cost of the dual-balancing policy associated with all periods $t \in \mathcal{T}_\Pi$ is at most the overall backlogging penalty incurred by OPT, denoted by Π^{OPT}, i.e., $\sum_{t \in \mathcal{T}_\Pi} \Pi_t^B \leq \Pi^{OPT}$ with probability* 1.

The next theorem follows as a corollary of Lemmas 1, 2, and 3.

Theorem 1 (Theorem 4.1 in Levi et al. [25]). *The dual-balancing policy for the uncapacitated periodic-review stochastic inventory control problem has a worst-case performance guarantee of 2, i.e., for each instance of the problem, the expected cost of the dual-balancing policy is at most twice the expected cost of an optimal solution, i.e., $\mathrm{E}[\mathcal{C}(B)] \leq 2\mathrm{E}[\mathcal{C}(OPT)]$.*

In Lemma 1, it was shown 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.,

$$\mathrm{E}[\mathcal{C}(B)] = 2 \sum_{t=1}^{T-L} \mathrm{E}[Z_t].$$

Lemmas 2 and 3 imply the following lower bound on the optimal expected cost of OPT. Using conditional expectations and the definition of Z_t, this implies that, indeed,

$$\begin{aligned}
\mathrm{E}[\mathcal{C}(OPT)] &\geq \mathrm{E}\Bigg[\sum_{t \in \mathcal{T}_H} H_t^B + \sum_{t \in \mathcal{T}_\Pi} \Pi_t^B\Bigg] \\
&= \sum_t \mathrm{E}\big[H_t^B \cdot \mathbb{1}(t \in \mathcal{T}_H) + \Pi_t^B \cdot \mathbb{1}(t \in \mathcal{T}_\Pi)\big] \\
&= \sum_t \mathrm{E}\big[\mathrm{E}\big[H_t^B \cdot \mathbb{1}(t \in \mathcal{T}_H) + \Pi_t^B \cdot \mathbb{1}(t \in \mathcal{T}_\Pi) \mid F_t\big]\big] \\
&= \sum_t \mathrm{E}\big[(\mathbb{1}(t \in \mathcal{T}_H) + \mathbb{1}(t \in \mathcal{T}_\Pi)) Z_t\big] = \sum_t \mathrm{E}[Z_t].
\end{aligned}$$

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 F_t, the quantities y_t^B and y_t^{OPT} are known deterministically, and so are the indicators $\mathbb{1}(t \in \mathcal{T}_H)$ and

$\mathbb{1}(t \in \mathcal{T}_\Pi)$. Suppose that the optimal policy is a randomized policy, i.e., it selects an order of size Q_t^{OPT} as a random function of f_t; then, the same arguments above still work, although now the conditioning is not only on F_t, but also on Q_t^{OPT}. Because the inventory control policy does not have any effect on the evolution of the future demands, the arguments above are still valid. This concludes the proof of the theorem. □

We note that the above worst-case analysis is tight in the sense that there exists a set of instances on which the ratio between the expected cost of the policy and the optimal expected cost converges to 2 (Levi et al. [25]). On the other hand, in computational experiments (Hurley et al. [13]) the dual-balancing policy perform significantly better than the worst-case guarantee, in many cases within a couple of percentages of optimum. Moreover, in many important settings the dual-balancing policies outperform the simple myopic policies.

2.2.4. The Capacitated Model. Next we briefly discuss the capacitated version of the model, in which there is potentially a capacity constraint $u_t < \infty$ on the size of the order placed in period t. (This follows closely Levi et al. [27].) Observe that a dual-balancing policy identical to the one in the uncapacitatetd case is not necessarily feasible because of the capacity constraints. (The balancing quantity q_s' might not be achievable.) The underlying reason is the fact that, unlike in the uncapacitated case, under ordering in a given period s might have implications beyond the single period $s+L$. Thus, one has to develop a *marginal backlogging cost accounting scheme.*

Consider some period s, and let again x_s denote the inventory position at the beginning of period s. Assume that the number of units ordered in the period is $q_s < u_s$, where again u_s is the maximum order size allowed in period 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 $t \ge s+L$ after this order arrives and becomes available. Suppose that for some realization $d_s, \dots, d_t$ of the demands $D_s, \dots, D_s$, respectively, we have that $d_{[s,\,t]} - (x_s + q_s + \sum_{j\in(s,\,t-L]} u_j) > 0$. The latter expression captures the backlogging cost of a policy that *after* period s always orders up to capacity. This implies that the shortage in period t and the resulting backlogging cost would exist, even if in every period over the time interval $[s, t-L]$ orders are placed up to the maximum available capacity. 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 $q_j \le u_j$ for each period j). In other words, given our decision not to order additional $\bar{q}_s$ in period s, the $d_{[s,\,t]} - (x_s + q_s + \sum_{j\in(s,\,t-L]} u_j) > 0$ part of the shortage could not be prevented by any decision made over the time 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 $\bar{q}_s$. In this case, we would say that this part of the backlogging cost in period t was *forced* by the decision in period s not to order up to capacity, 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 $t-L$.

Let W_{st} be the shortage in period t that is forced by the decision in period s. 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., $\tilde{\Pi}_s^P = \sum_{t=s+L}^{T} p_t W_{st}^P$. (See Figure 1 for a graphical illustration.) Similar to the uncapacitated case, express the cost of each policy P as $\mathcal{C}(P) = \sum_{t=1}^{T-L} (H_t^P + \tilde{\Pi}_t^P)$. The function $\pi_s^B(q_s) := \mathrm{E}[\tilde{\Pi}_s^B(q_s) \mid f_s]$ is a continuous, nonincreasing function of q_s that is equal to 0 for $q_s = u_s$. Thus, the dual-balancing policy can be implemented in a manner similar to the uncapacitated case. Moreover, Levi et al. [27] have shown that the policy has a worst-case performance guarantee of 2.

Figure 1. Allocation of a period backorder to ordering decisions (Levi et al. [27]).

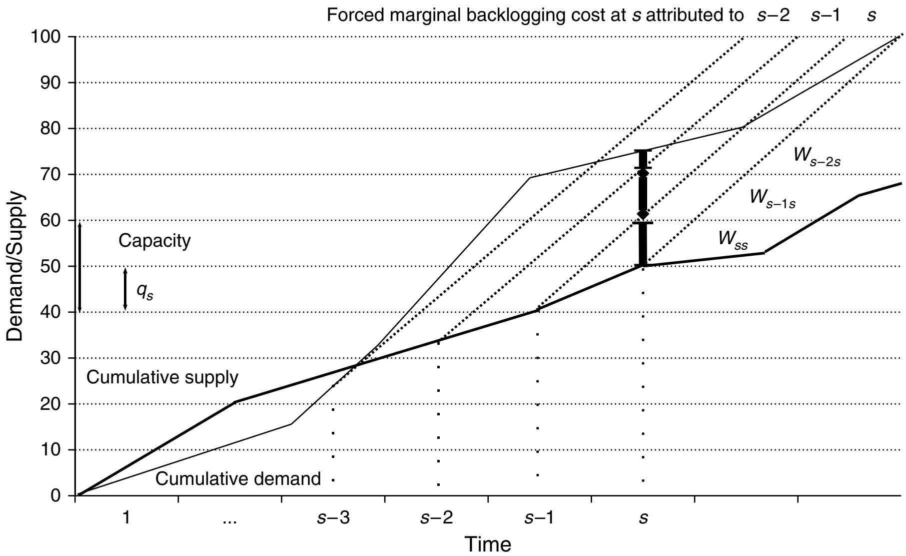

2.2.5. Concluding Comments and Extensions. In a stream of follow-up work to Levi et al. [27] and [25], the ideas of marginal cost accounting schemes and cost-balancing techniques have been used to devise provably near-optimal approximation algorithms for a large class of core stochastic inventory control models. In particular, Levi and Shi [20] study stochastic lot-sizing models with positive fixed ordering cost, Levi et al. [21] study models with lost sales, and Levi et al. [23] study multiechelon models. In addition, it has been shown that one can consider parametric versions of these policies and use known lower and upper bounds on the optimal ordering levels to devise more sophisticated policies with better empirical (typical) performance, some of which still admit worst-case analysis. The parametrization of these policies is motivated by the fact that, in the context of worst-case analysis, one wishes to choose policies that "protect" against all possible instances, whereas per a given (known) instance it is possible to choose the parameters of the policy optimally with respect to that instance. This naturally leads to improved empirical performance. Moreover, computational experiments indicate that cost-balancing-based algorithms typically perform significantly better than what the respective worst-case analysis suggests, and in many important scenarios significantly outperform previously known heuristics (Levi et al. [27]; Hurley et al. [13]; Levi and Shi [20]), and can be applied in very general settings.

The extension of these ideas to additional core models in stochastic inventory theory, such as the *one-warehouse multiretailer problem* is one of the major current challenges in this area of research. Another important future research direction is to study the performance of dual-balancing or, more generally, cost-balancing policies under various assumptions on the underlying demand distributions. As much as it is powerful to establish general worst-case analysis, it is equally important to refine this analysis to various parametric regimes of the underlying demand distributions and other key parameters of the problem. We call this *parametric worst-case analysis*.

Another important direction is to establish a better rigorous understanding of the computational complexity of various stochastic inventory control models. There are relatively few results on the computational complexity of these models (Halman et al. [11]), and the analysis seems to require more advanced techniques from complexity theory of algorithms.

Finally, it would be interesting to apply these new algorithmic ideas in other application domains to obtain provably near-optimal algorithms with worst-case guarantees.

3. Revenue Management of Reusable Resources: Linear-Programming-Based Policies

This section follows closely Levi and Radovanovic [19], who studied the revenue management of systems with reusable resources. These models capture several core applications in revenue management, such as workforce management, hotel room allocation, and car rental management. Like the previous stochastic inventory models, these problems can be formulated as dynamic programs with typically large state space that does not lead to algorithms with reasonable computational tractability. In contrast, we shall discuss how to use a simple *linear program* (LP) to generate a relaxation of the original problem, and then leverage the optimal solution of the LP to construct a simple stochastic control policy. The performance analysis is based on well-known *loss network* models from stochastic analysis (Kelly [17]). In particular, we show that the proposed policy has a worst-case guarantee of 2 and then use parametric worst-case analysis to show that the policy is in fact asymptotically optimal with respect to a certain parametric regime.

3.1. Model Formulation

Consider a *revenue management model with a single reusable resource*, where a single resource is used to serve multiple classes of customers. Specifically, there is a single resource pool of integer capacity $C < \infty$ that is used to satisfy the demands of M different classes of customers. The customers of each class $i = 1, \ldots, M$ arrive according to an independent Poisson process with respective rate λ_i. Each class i customer requests $A \in \mathbb{Z}_+$ units of the resource (without loss of generality, $A = 1$) for a certain period of time that is a priori random following a general class-specific distribution with mean μ_i that is independent of the arrival process and between different customers. During the time a customer is served, the requested unit can not be used by any other customer; after the service is over, the unit becomes available again to serve other customers. (While a customer is served, only the conditional distribution of the residual service time of this customer is known.) Class i customers pay p_i dollars per unit of service time. Customers can be served only if upon their arrival there is at least 1 unit of available capacity. However, customers can be rejected even if there is available capacity. (Rejecting a customer now possibly enables serving more profitable customers in the future.) Customers that are not served upon arrival are *lost* and leave the system. The goal is to find an admission policy that maximizes the expected long-run revenue rate. Specifically, if $\mathcal{R}_\pi(T)$ is the revenue achieved by policy π over the interval $[0, T]$, then the expected long-run revenue rate of π, is defined as $\mathcal{R}(\pi) =: \liminf_{T\to\infty}(\mathbb{E}_\pi[\mathcal{R}_\pi(T)]/T)$, where the expectation $\mathbb{E}_\pi$ is taken with respect to the probability measure induced by fixing policy π.

Like many stochastic optimization models, one can formulate this problem using a dynamic programming approach. However, even in special cases (e.g., with class-specific exponentially distributed service times), the resulting dynamic program seems computationally intractable because the corresponding state space grows very fast.

3.2. An LP-Based Approach

Next we construct a simple LP that provides an upper bound on the achievable long-run average revenue rate. The LP is identical to the one used by Key [18], Hunt and Laws [12], and Iyengar and Sigman [14], and is also similar to the one used by Adelman [1] in the queueing networks framework with unit resource requirements. Several LP-based policies have been proposed to the model discussed above (Key [18], Kelly [17], Iyengar and Sigman [14]). However, their performance has been analyzed in a very specific *heavy traffic regime*. In contrast, we shall show how to use the optimal solution of the LP to construct a simple admission control policy that is called *class selection policy* (CSP). Moreover, the LP will be used to analyze the performance of the proposed policy and provide both *uniform* worst-case guarantees, parametric worst-case analysis, as well as asymptotic analysis in a parametric regime that is more general.

It is well known that state-dependent policies are optimal for the model considered above. Moreover, any state dependent policy induces a Markov process on the state space of the system with a unique stationary distribution that is ergodic. (The state of the system is the number of customers of each class currently being served and the time elapsed from the moment their service has started.) It follows that each given state-dependent policy π induces stationary probabilities $\alpha_i^{(\pi)}$ for each $i = 1, \ldots, M$ of accepting a class i customer in some random time. (This captures the long-run proportion of class i customers being served under policy π.) Furthermore, by applying Little's law, we can use the stationary probability $\alpha_i^{(\pi)}$ to express the long-run expected number of class i customers being served in the system under the state-dependent policy π as $\lambda_i \mu_i \alpha_i^{(\pi)} = \rho_i \alpha_i^{(\pi)}$. (Note that $\rho_i = \lambda_i \mu_i$ is the expected number of class i customers being served in the system with infinite capacity and no rejections; in the context of communication networks it is usually called the *traffic intensity*.) It follows that the expected long-run revenue rate of policy π can be expressed as $\sum_{i=1}^{M} p_i \alpha_i^{(\pi)} \lambda_i \mu_i = \sum_{i=1}^{M} p_i \alpha_i^{(\pi)} \rho_i$. Similarly, the overall expected long-run number of customers being served can be expressed as $\sum_{i=1}^{M} \alpha_i^{(\pi)} \rho_i$. Moreover, it is clear that for every feasible policy the number of customers being served is limited by the capacity. In particular, this implies that $\sum_{i=1}^{M} \alpha_i^{(\pi)} \rho_i \leq C$. This gives rise to the following LP:

$$\max_{\alpha_1, \ldots, \alpha_M} \quad \sum_{i=1}^{M} p_i \alpha_i \rho_i \tag{6}$$

$$\text{s.t.} \quad \sum_{i=1}^{M} \alpha_i \rho_i \leq C, \tag{7}$$

$$0 \leq \alpha_i \leq 1 \quad \forall 1 \leq i \leq M. \tag{8}$$

Note that, for each feasible state-dependent policy π, the vector $\alpha^{(\Pi)} = (\alpha_1^{(\pi)}, \alpha_2^{(\pi)}, \ldots, \alpha_M^{(\pi)})$ is a feasible solution for the LP defined by (6)–(8) above and has objective value that is equal to the expected long-run revenue rate of policy π. In fact, the LP enforces the capacity constraint of the system only in expectation, whereas in the original problem this constraint has to hold for every sample path. It follows that the LP defined by (6)–(8) relaxes the original problem and provides an upper bound on the optimal (best achievable) expected long-run revenue rate.

Moreover, the LP defined above is a *knapsack* LP. In particular, the optimal solution has the following structure: $\alpha_1 = \alpha_2 = \cdots = \alpha_{M'-1} = 1$, for some M' (where $1 \leq M' \leq M$); the value of $\alpha_{M'}$ is possibly a fraction, i.e., $0 < \alpha_{M'} \leq 1$; and, for $i = M' + 1, \ldots, M$, we have $\alpha_i = 0$.

3.2.1. The Class Selection Policy. Let $\alpha^* = (\alpha_1^*, \ldots, \alpha_M^*)$ be the optimal solution of the knapsack LP defined by (6)–(8) above. Consider the following simple policy that is called the *class selection policy*. For each $i = 1, \ldots, M' - 1$, *accept* each arriving customer as long as there is available capacity in the system upon arrival. If $i = M'$, *accept* with probability $0 < \alpha_{M'} \leq 1$ and as long as there is available capacity upon arrival. For each $i = M' + 1, \ldots, M$, *reject*.

The CSP has a very simple structure. It always admits customers from the classes for which the corresponding value α_i^* in the optimal LP solution equals 1 as long as capacity permits. It never admits customers from classes for which the corresponding value α_i^* equals 0, and it flips a coin for the possibly one class with fractional value $\alpha_{M'}^*$. The CSP is conceptually very intuitive in that it splits the classes into profitable and nonprofitable ones that should be ignored.

3.2.2. Performance Analysis. For simplicity of exposition, assume that α^* is integral (i.e., $\alpha_i^* \in \{0,1\}$, for each $i = 1,\ldots,M$). The CSP induces a well-structured stochastic process. Specifically, Each class $i = 1,\ldots,M$ induces a Poisson arrival stream with respective rate $\alpha_i^*\lambda_i$, $1 \le i \le M$. Thus, for each class i with $\alpha_i^* = 1$, the arrival process is identical to the original process and each class i with $\alpha_i^* = 0$ can be ignored. This gives rise to system dynamics analogous to the classical *loss network model* with a single resource. There are C servers (units) that are used to serve M' independent Poisson streams of requests. The requests of stream (class) i arrive at rate λ_i, and each requires 1 server (unit) for some random service time with mean μ_i. A request is served if upon arrival there is an idle server (available unit); otherwise, it is rejected and leaves the system (lost). Loss networks have been studied extensively in the context of communication networks, and there is a huge body of literature on this topic (see the literature review in Levi and Radovanovic [19]).

Next, we are interested in characterizing the long-run *blocking probability* of class i customers induced by the CSP. The blocking probability of a class i customer is the stationary probability that the customer arrives at a random time to the system and is rejected by the CSP because there is no available capacity. For each $i = 1,\ldots,M'$, let Q_i be the stationary probability of rejecting a class i customer under the CSP. To express Q_i, we let X_i be the stationary (random) number of class i customers being served in the system at some random time under the CSP. By the PASTA (poisson arrivals see time averages) property (see Wolff [41]), it follows that $Q_i = Q = \Pr(\sum_{k=1}^{M'} X_k = C)$, for each $i = 1,\ldots,M'$. Moreover, because the corresponding stochastic process is ergodic, it follows that the long-run proportion of class i customers being served is equal to $1 - Q_i = 1 - Q$. Thus, the expected long-run revenue rate of the CSP can be expressed as $\sum_{i=1}^{M'} p_i\alpha_i^*\rho_i(1 - Q_i) = \sum_{i=1}^{M'} p_i\rho_i(1 - Q)$. However, $\sum_{i=1}^{M'} p_i\alpha_i^*\rho_i$ is the optimal value of the LP, which is an upper bound on the best achievable expected long-run revenue rate, denoted by $\mathcal{R}(OPT)$. Thus, a key aspect of the performance analysis of the CSP is to obtain a lower bound on the probability $1 - Q$ or, equivalently, an upper bound on the probability Q. Specifically, if $1 - Q \ge \beta$, it follows that $\mathcal{R}(CSP) = \sum_{i=1}^{M'} p_i\alpha_i^*\rho_i(1 - Q) \ge \sum_{i=1}^{M'} p_i\alpha_i^*\rho_i\beta \ge \beta\mathcal{R}(OPT)$; that is, the CSP is a β-approximation.

Levi and Radovanovic [19] obtained an upper bound on the blocking probability Q.

Lemma 4 (Lemma 1 in Levi and Radovanovic [19]). *For each $i = 1,\ldots,M'$, the blocking probability Q_iQ has the following properties:*

(i) $Q \le 0.5$ *for all capacity values* $C \ge 1$.

(ii) *As C grows to infinity, the probability Q diminishes to* 0 *at rate of at least* $O(1/\sqrt{C})$ *regardless of other parameters of the problem such as the price rates, service time distributions, and number of classes.*

Next we briefly discuss the proof of Lemma 4. Using the results from Sevastyanov [33], Kelly [16], Kaufman [15], Burman et al. [4], and Whitt [40], which characterize the stationary distributions of the corresponding loss network models, one can develop explicit (exact) expressions for the resulting blocking probability Q induced by the CSP. The expression has the form of Erlang formula and can be upper bounded by an expression that depends *only* the capacity of the system C, regardless of the other parameters such as arrival rates, number of classes, prices, and service time distributions, which can be general, not necessarily exponential. The latter expression can then be bounded uniformly by 0.5 (see Part (i) of Lemma 4 above) for each possible value of capacity $C \ge 1$. In addition, it can be shown that as the capacity of the system goes to infinity, the upper bound goes to 0, namely, the blocking probability Q goes to 0 (Part (ii) of Lemma 4). This implies that the CSP becomes optimal in this asymptotic regime. Lemma 4 implies that the CSP is guaranteed to have an expected long-run revenue rate that is at least half of optimal for *all* possible values of the resource capacity C regardless of the prices, arrival rates, and service time distributions. Moreover, the bound becomes better as the capacity of the system grows large (whereas all the other parameters can change arbitrarily), and the CSP is asymptotically optimal. In contrast, previous asymptotic optimality results are obtained only in very specific regimes, usually called heavy traffic regimes. The capacity and the arrival rates are scaled simultaneously at the same (linear) rate, whereas all the other parameters of the problem, such as service time distributions, resource requirements, number of classes, and price rates, are kept fixed. Some of the results also require the assumption that the service times are exponentially distributed. (We note that a similar notion of this heavy traffic regime has been used to analyze the asymptotic performance of other heuristics in other applications, e.g., Gallego and van Ryzin [10].) Although these might be reasonable assumptions in the context of communication networks, they are less likely to hold in the other application domains.

3.2.3. Extensions and Future Directions. The CSP and its analysis can be extended to a model with *static pricing*. In this model, one first has to determine the respective prices that each class is charged. The respective arrival rate of each class depends on the price it is charged. After the prices are set, one seeks to find the best admission control policy that maximizes the expected long-run revenue rate. However, the policy in this more general model is derived based on a *nonlinear program* (NLP). Levi and Radovanovic [19] show how to simplify the resulting NLP, and in what scenarios it can be solved efficiently.

In many applications (e.g., hotel rooms or car rentals), customers do not necessarily start the service upon arrival, but rather place a reservation in advance for some interval of time in the future. These models are significantly more complex and have the structure of a *loss network with advance reservations*. There is relatively little work on these more complex systems, and the only existing results are for very special and restrictive cases.

Linear programming and, more generally, mathematical programming have been used extensively to obtain relaxations and provably near-optimal approximation algorithms for deterministic combinatorial optimization problems (Vazirani [39]). More recent work has extended this approach to two-stage stochastic optimization models (Shmoys and Swamy [34, 35]) and multistage stochastic optimization problems (see, for example, Möhring et al. [29, 30], Dean et al. [8]). This seems like a promising direction that could lead to additional important algorithmic results for multistage optimization models, including core models in operations management.

Acknowledgments

The author thanks Mahesh Nagarajan, Georgia Perakis, and David Shmoys for reading and commenting on a draft of this survey. This work was supported in part by National Science Foundation Grants DMS-0732175 and CMMI-0846554 (CAREER Award), Air Force Office of Scientific Research Award FA9550-08-1-0369, an Singapore MIT Alliance grant, and the Buschbaum Research Fund of MIT.

References

[1] D. Adelman. Price-directed control of a closed logistics queueing network. *Operations Research* 55:1022–1038, 2007.

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

[3] S. Axsäter and P. Lundell. In process safety stock. *Proceedings of the 23rd IEEE Conference on Decision and Control*, IEEE Control Systems Society, Las Vegas, NV, 839–842, 1984.

[4] D. Y. Burman, J. P. Lehoczky, and Y. Lim. Insensitivity of blocking probabilities in a circuit-switching network. *Journal of Applied Probability* 21:850–859, 1984.

[5] C. W. Chan and V. F. Farias. Stochastic depletion problems: Effective myopic policies for a class of dynamic optimization problems. *Mathematics of Operations Research* 34:333–350, 2009.

[6] F. Chen. 94%-effective policies for a two-stage serial inventory system with stochastic demand. *Management Science* 45(12):1679–1696, 1999.

[7] L. Y. Chu and Z. J. M. Shen. A power-of-two ordering policy for one-warehouse multiretailer systems with stochastic demand. *Operations Research* 58:492–502, 2010.

[8] B. C. Dean, M. X. Goemans, and J. Vondrák. Approximating the stochastic knapsack problem: The benefit of adaptivity. *Proceedings of the 45th Annual IEEE Symposium on the Foundations of Computer Science*, IEEE Computer Society, Washington, DC, 208–217, 2004.

[9] S. Dye, L. Stougie, and A. Tomasgard. The stochastic single resource service provision problem. *Naval Research Logistics* 50:869–887, 2003.

[10] G. Gallego and G. van Ryzin. Optimal dynamic pricing of inventories with stochastic demands over finite horizons. *Management Science* 40:999–1020, 1994.

[11] N. Halman, D. Klabjan, M. Mostagir, J. Orlin, and D. Simchi-Levi. A fully polynomial-time approximation scheme for single-item stochastic inventory control with discrete demand. *Mathematics of Operations Research* 34:674–685, 2009.

[12] P. J. Hunt and C. N. Laws. Optimization via trunk reservation in single resource loss systems under heavy traffic. *Annals of Applied Probability* 7:1058–1079, 1997.

[13] G. Hurley, P. Jackson, R. Levi, R. O. Roundy, and D. B. Shmoys. New policies for stochastic inventory control models: Theoretical and computational results. Working paper, Massachusetts Institute of Technology, Cambridge, http://www.retsef.scripts.mit.edu/, 2005.

[14] G. Iyengar and K. Sigman. Exponential penalty function control of loss networks. *Annals of Applied Probability* 14:1698–1740, 2004.

[15] J. S. Kaufman. Blocking in a shared resources environment. *IEEE Transactions on Communications* 29:1474–1481, 1981.

[16] F. P. Kelly. *Reversibility and Stochastic Networks*. John Wiley & Sons, Inc., New York, 1979.

[17] F. P. Kelly. Effective bandwidths at multi-class queues. *Queueing Systems* 9:5–16, 1991.

[18] P. Key. Optimal control and trunk reservation in loss networks. *Probability in Engineering and Informations Sciences* 4:203–242, 1990.

[19] R. Levi and A. Radovanovic. Provably near-optimal LP-based policies for revenue management in systems with reusable resources. *Operations Research* 58(2):503–507, 2010.

[20] R. Levi and C. Shi. Approximation algorithms for the stochastic lot-sizing problem. Working paper, Massachusetts Institute of Technology, Cambridge, MA, 2009.

[21] R. Levi, G. Janakiraman, and M. Nagarajan. A 2-approximation algorithm for stochastic inventory control models with lost sales. *Mathematics of Operations Research* 33:351–374, 2008.

[22] R. Levi, A. Lodi, and M. Sviridenko. Approximation algorithms for the multi-item capacitated lot-sizing problem via flow-cover inequalities. *Mathematics of Operations Research* 33(2):461–474, 2008.

[23] R. Levi, R. O. Roundy, and V. A. Truong. Provably near-optimal balancing policies for multi-Echelon stochastic inventory control models. Working paper, Massachusetts Institute of Technology, Cambridge, MA, 2005.

[24] R. Levi, R. O. Roundy, and D. B. Shmoys. Primal-dual algorithms for deterministic inventory problems. *Mathematics of Operations Research* 31:267–284, 2006.

[25] R. Levi, M. Pál, R. O. Roundy, and D. B. Shmoys. Approximation algorithms for stochastic inventory control models. *Mathematics of Operations Research* 32:284–302, 2007.

[26] R. Levi, R. O. Roundy, D. B. Shmoys, and M. Sviridenko. A constant approximation algorithm for the one-warehouse multi-retailer problem. *Management Sceince* 54:763–776, 2008.

[27] R. Levi, R. O. Roundy, D. B. Shmoys, and V. A. Truong. Approximation algorithms for capacitated stochastic inventory control models. *Operations Research* 56:1184–1199, 2008.

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

[29] R. H. Möhring, F. J. Radermacher, and G. Weiss. Stochastic scheduling problems I: General strategies. *Zeitschrift für Operations Research* 28:193–260, 1984.

[30] R. H. Möhring, F. J. Radermacher, and G. Weiss. Stochastic scheduling problems II: Set strategies. *Zeitschrift für Operations Research* 29:65–104, 1984.

[31] R. H. Möhring, A. Schulz, and M. Uetz. Approximation in stochastic scheduling: The power of LP-based priority policies. *Journal of the ACM* 46:924–942, 1999.

[32] R. O. Roundy. Efficient, effective lot-sizing for multi-product, multi-stage production systems. *Operations Research* 41:371–386, 1993.

[33] B. A. Sevastyanov. An ergodic theorem for Markov processes and its application to telephone systems with refusals. *Theory of Probability and Its Applications* 2:104–112, 1957.

[34] D. B. Shmoys and C. Swamy. Stochastic optimization is (almost) as easy as deterministic optimization. *Proceedings of the 45th Annual IEEE Symposium on the Foundations of Computer Science*, IEEE Computer Society, Washington, DC, 228–237, 2004.

[35] D. B Shmoys and C. Swamy. Approximation algorithms for 2-stage stochastic optimization problems. *ACM SIGACT News Archive* 37:33–46, 2006.

[36] E. A. Silver and H. C. Meal. A heuristic selecting lot-size requirements for the case of a deterministic time varying demand rate and discrete opportunities for replenishment. *Production and Inventory Management* 14:64–74, 1973.

[37] L. Stougie and M. H. van der Vlerk. Approximation in stochastic integer programming. SOM Research Report 03A14, Eindhoven University of Technology, Eindhoven, The Netherlands, 2003.

[38] C. Swamy and D. B. Shmoys. Sampling-based approximation algorithms for multi-stage stochastic optimization. *Proceedings of the 46th Annual IEEE Symposium in the Foundations of Computer Science,* IEEE Computer Society, Washington, DC, 357–366, 2005.

[39] V. J. Vazirani. *Approximation Algorithms.* Springer, Berlin, 2001.

[40] W. Whitt. Blocking when service is required from several facilities simultaniously. *AT&T Technical Journal* 64:1807–1856, 1985.

[41] R. W. Wolff. *Stochastic Modeling and Theory of Queues.* Prentice Hall, Berlin, 1989.

[42] P. H. Zipkin. *Foundations of Inventory Management.* McGraw-Hill, Boston, 2000.