

A Fully Polynomial Time Approximation Scheme for Single-Item Stochastic Inventory Control with Discrete Demand

N. Halman* D. Klabjan* M. Mostagir* J. Orlin † D. Simchi-Levi*‡

Abstract

The single-item stochastic inventory control problem is to find an inventory replenishment policy in the presence of independent discrete stochastic demands under periodic review and finite time horizon. In this paper, we prove that this problem is intractable and for any $\epsilon > 0$, we design an algorithm with running time polynomial in the size of the problem input and in $1/\epsilon$, that finds a policy whose value is within a factor $(1 + \epsilon)$ of the value of an optimal policy. In addition, we formally prove that finding an optimal policy is intractable.

1 Introduction

The standard single-item stochastic inventory control problem is to find replenishment quantities in each time period that minimize the expected procurement and holding/backlogging cost. We assume dynamic time replenishment over a finite number of time periods. We assume also that the holding cost, which includes a potential penalty for backlogging, is convex. In addition, the procurement cost is convex and nondecreasing. In the case of linear procurement costs, it is well known that the base stock policy is optimal, see e.g., [SCB05, Zip00]. This policy assumes a time-dependent number called the base-stock level, and the policy places an order that brings the inventory level up to the base-stock level. If the inventory level is above the base-stock level, then no order is placed. While this theory shows the existence of base-stock levels, it does not provide an efficient way to compute these levels.

In the linear procurement cost case, the computation of the optimal base-stock levels is a nontrivial task. It is for this reason that we resort to approximation algorithms. Our main result is a fully polynomial time approximation scheme (FPTAS) for the aforementioned stochastic inventory control problem. A minimization problem has an FPTAS if for every $\epsilon > 0$ and for every instance \mathcal{I} we have

$$\mathcal{A}(\mathcal{I}) \leq (1 + \epsilon)\text{opt}(\mathcal{I}), \tag{1}$$

where $\text{opt}(\mathcal{I})$ is the optimal value and $\mathcal{A}(\mathcal{I})$ is the value returned by the approximation algorithm \mathcal{A} . The running time of algorithm \mathcal{A} must be polynomial in $1/\epsilon$ and the size of the problem input.

We assume that the demands are discrete, independent random variables, but not necessarily identically distributed. We also assume that the distributions are known in advance and there is zero lead time. Although the policy structure in the general convex nondecreasing procurement cost case is not known, our FPTAS works also in this case.

The standard dynamic programming approach gives a pseudo-polynomial algorithm. This algorithm is linear with respect to the maximum demand value, which is exponential in the input size. This dynamic

*Research supported by ONR Contracts N00014-95-1-0232 and N00014-01-1-0146 and NSF Contracts DMI-0085683 and DMI-0245352, and NASA interplanetary supply chain and logistics architectures project.

†Research supported in part by ONR grant N00014-98-1-0317.

N. Halman, M. Mostagir, J. Orlin and D. Simchi-Levi: Massachusetts Institute of Technology, Cambridge, MA, E-mail: {halman,mostagir,jorlin,dslevi}@mit.edu

D. Klabjan: Northwestern University, Evanston, IL, E-mail: d-klabjan@northwestern.edu

‡Research is also supported by NSF Contracts CMMI-0758069

program considers an exponential number of different inventory levels. To circumvent this, in each time period we carefully select only a subset of possible inventory levels and compute an approximation of the optimal value function on this subset. At all other inventory levels, the approximate value function is interpolated. Our approach differs from other FPTAS's, that rely on scaling and rounding the data to reduce the state space.

We note that Woeginger's general framework for transforming a DP into an FPTAS [Woe00] does not include our problem. First of all, his approach does not consider cases in which the action space is exponentially large. Second, he does not consider applications to stochastic optimization problems.

Perhaps the most closely related technique methodologically to our approach is that of [DGV08], who also rely on using an approximation function combined with linear interpolation (see Section 6, proof of Lemma 9 in their work). In addition to the standard inventory control problem, we also show how to obtain an FPTAS from our framework for many extensions: the capacitated and discounted versions, the lost sales model, the model where disposal of excess inventory at a cost is allowed, and a model where only a non-exact evaluation of the cost functions is available. We also provide a reduction that shows the NP-hardness of the stochastic problem with linear procurement and holding costs.

The main methodological contribution of this work is the development of FPTASs using approximate, efficiently-representable value functions. While this concept is widely used in computational dynamic programming algorithms, to the best of our knowledge, we are the first to use this approach to develop FPTASs. Another key contribution is the hardness proof of the basic stochastic inventory control problem.

The manuscript is structured as follows. We introduce and state the stochastic inventory control problem in Section 3. In Section 4, we show that this problem is NP-hard. The framework of approximation sets and functions is given in Section 5. In Section 6, we present the FPTAS and its analysis. In Section 7, we provide FPTASs for several extensions of the basic inventory control problem. We conclude the introduction with a literature review.

Literature review

The stochastic inventory control problem is one of the most widely studied problems in inventory theory. The so called (s, S) policy, which results from the economies of scale in procurement, is a widely used policy in practice. We refer the reader to the books [SCB05, Por02, Zip00] for an in-depth coverage of the topic.

For the finite time horizon inventory control problems with linear cost functions, the base-stock levels can be computed in pseudopolynomial time recursively from the optimality equations. For infinite horizon problems, the optimality equation still holds. [ZF91] and [FZ84b] propose algorithms for finding an optimal policy for the infinite horizon inventory control problem. The multi-echelon problem, [FZ84a], and a variant of inventory control with fixed cost for backlogging, [RM00], have also been considered, although no efficient approximation algorithms have been developed for these problems.

The deterministic capacitated lot-sizing problem was solved with FPTASs by [VW01] and [SO95]. However, their approach does not appear to extend to the stochastic lot-sizing problem.

Recently there is a growing interest in approximation algorithms for stochastic problems. [SS06] provide an FPTAS for a broad class of linear 2-stage stochastic recourse problems, even in the case of exponentially many second stage variables. Their algorithm is based on the standard convex programming formulation. They solve the linear programming relaxation by the ellipsoid algorithm and then they round the solution to the closest integral vector. The multi-stage setting is considered in [SS05].

A 2-approximation algorithm for stochastic inventory control is presented in [LPRS07]. They assume stochastic lead times with no order crossing, possibly correlated demand distributions, and the "no speculative cost assumption", which is equivalent to assuming that there is no ordering cost. The first two assumptions are more general than our assumptions, whereas their latter assumption is more restrictive. They study the so called dual-balancing policy, which is not a base-stock policy. An extension to the capacitated version is given in [LRST08]. In [LRS07], the authors use sampling to analyze the inventory control problem with convex cost functions. Their model is more general than ours since they do not assume an explicit knowledge of the demand distribution. The algorithm is based on sampling, and they provide approximate base-stock levels. The number of required samples to achieve an ϵ approximation is upper

bounded by $1/\epsilon$, the number of time periods, and $(h + b)^2/\min\{b, h\}^2$, where h is the stationary per unit holding cost and b is the stationary per unit backlog penalty. Note that this bound is not polynomial in the input size and therefore their algorithm is possibly exponential in the size of the instance and in $\frac{1}{\epsilon}$. In addition, their algorithm is not deterministic.

The approaches in [SS06] and [LRS07] are significantly different. While both use sampling, the former relies heavily on optimization techniques and the latter exploits the underlying nature of stochastic inventory problems. To the contrary of these two approaches, we do not use sampling and our approach is more combinatorial in nature.

2 Notation

Let $\mathbb{R}, \mathbb{Z}, \mathbb{Z}_+, \mathbb{N}$ denote the set of real numbers, integers, nonnegative integers, and positive integers, respectively. For any pair A, B of integers with $-\infty < A < B < \infty$, let $[A, B] = \{A, A+1, \dots, B\}$ denote the set of integers between A and B . For any function f over $[A, B]$ we denote by f^{\max} the maximal value f achieves over $[A, B]$. Let $U = B - A + 1$ be the number of points in the domain, and $\bar{U} = f^{\max}$ be the maximal value f achieves on the domain. For any subset $D' \subseteq [A, B]$, a function f is called *unimodal over D'* if there exist $x^* \in D'$ such that f is nonincreasing over $[A, x^*]$ and nondecreasing over $[x^*, B]$. The value x^* is called an arg min of f . For any subset $D' \subseteq [A, B]$, we define the *piecewise linear extension of f induced by D'* as the continuous function obtained by making f linear between successive values of D' . For any subset $D' \subseteq [A, B]$, a function f on $[A, B]$ is called *convex over D'* if its piecewise linear extension induced by D' is convex. For any function f on $[A, B]$ and subset $D' \subseteq [A, B]$, we define the *convex extension of f induced by D'* to be the piecewise linear extension of f induced by $D'' \subseteq D'$, where D'' is the projection of the convex hull of $\{(x, f(x)) \mid x \in D'\}$ on the x -axis. The base two logarithm of z is denoted by $\log z$.

3 Problem statement

Let T be the length of the planning horizon. At the beginning of a time period, we observe the inventory level and a replenishment decision is made. If we place an order, it arrives immediately, i.e., we assume there is no lead time. Backlogging is represented as a negative inventory level. A single convex holding cost function is used to model both backlogging and positive inventory. The holding cost is accounted for at the end of the time period. For each time period $t = 1, \dots, T$ we define:

- x_t : procurement quantity in time period t ;
- I_t : inventory level at the beginning of time period t ;
- \bar{I}_t : inventory level at the end of period t (i.e., $I_t = \bar{I}_{t-1}$);
- $c_t(x_t)$: procurement cost in time period t , given an order of size x_t ;
- $h_t(\bar{I}_t)$: holding cost in time period t , given inventory level \bar{I}_t .

For each time period $t = 1, \dots, T$, we assume that there is an oracle that computes functions c_t, h_t , and there is a discrete random variable D_t describing the demand in time period t . For each D_t , we are given n_t , the number of different values it admits with positive probability, and the demand values $d_{t,1} < \dots < d_{t,n_t}$. Moreover, we are also given positive integers $q_{t,1}, \dots, q_{t,n_t}$ such that

$$\text{Prob}[D_t = d_{t,i}] = \frac{q_{t,i}}{\sum_{j=1}^{n_t} q_{t,j}}.$$

The random variables D_t are assumed to be independent for different t . We define for every $t = 1, \dots, T$ and $i = 1, \dots, n_t$ the following values:

$$\begin{aligned}
p_{t,i} &= \text{Prob}[D_t = d_{t,i}] && \text{probability that there is a demand of } d_{t,i} \text{ units in time period } t; \\
n^* &= \max_t n_t && \text{maximum number of different values } D_t \text{ can take over the entire} \\
&&& \text{time horizon;} \\
d^* &= \max_t d_{t,n_t} && \text{maximum demand over the entire time horizon;} \\
D^* &= \sum_{t=1}^T d_{t,n_t} && \text{maximum total demand over the entire time horizon;} \\
Q_t &= \sum_{j=1}^{n_t} q_{t,j} && \text{a common denominator of all the probabilities in time period } t; \\
M_t &= \prod_{j=t}^T Q_j && \text{a common denominator of all the probabilities in all time} \\
&&& \text{periods following time period } t - 1; \\
M_{T+1} &= 1.
\end{aligned}$$

We make the following assumptions.

Assumption 3.1. *All demand, procurement and inventory levels are integral. Moreover, the demand and procurement levels are nonnegative.*

Assumption 3.2. *The procurement cost function c_t is nondecreasing nonnegative convex over \mathbb{Z}_+ for every $t = 1, \dots, T$.*

Assumption 3.3. *The holding cost function h_t is nonnegative convex over \mathbb{Z} for every $t = 1, \dots, T$.*

Assumption 3.4. *All cost functions can be evaluated in polynomial time at any value in their domain.*

Note that the binary input size of the problem is bounded below by

$$\Omega(T + n^* + \log d^* + \log \max_t \{c_t(D^*), h_t(D^*), h_t(-D^*)\}).$$

The last term takes into account the space needed for storing the value of c_t or h_t at D^* (see more details about this aspect in Section 6.1).

The objective is to minimize the total expected cost. The problem can be formulated as finding a policy $x_t = x_t(I_t)$ for $t = 1, \dots, T$ that realizes

$$z^* = \min_{x_t} E_D \left(\sum_{t=1}^T c_t(x_t) + h_t(I_t + x_t - D_t) \right),$$

subject to the system dynamics

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

The action space requirement is $x_t \in \mathbb{Z}_+$ for $t = 1, \dots, T$, and the initial state is $I_1 = 0$.

Note that in the above context we have $E_D(\sum_{t=1}^T c_t(x_t) + h_t(I_t + x_t - D_t)) = \sum_{t=1}^T c_t(x_t) + \sum_{j=1}^{n_t} p_{t,j} h_t(I_t + x_t - d_{t,j})$. The standard optimality equation is presented later in Section 6.1.

4 Hardness result

In this section we show that the single-item stochastic inventory control problem is #P-hard (and thus NP-hard) even in the case of linear procurement and holding costs. We make a transformation from a problem concerning the evaluation of convolutions of discrete variables, which in turn is proved #P-hard from a transformation from the K th largest subset problem (see problem SP20 on page 225 of [GJ79] for its NP-hardness and the same reference for its #P-hardness).

Problem: K th largest subset

Instance: A finite set $A = \{a_1, \dots, a_n\}$ of nonnegative integers and positive integer numbers $K \leq 2^n$ and B .

Question: Are there K or more distinct subsets $A' \subseteq A$ for which $\sum_{a \in A'} a \leq B$?

Remark: This problem is named the “ K th largest subset problem”, but it actually looks for the K th smallest subset of A .

In the counting version of the problem, one is asked to count the number of distinct subsets of A whose sum is at most B .

Problem: Evaluating the cdf of convolutions of discrete random variables

Instance: *Discrete random variables X_1, \dots, X_n and probabilities $p_{i,j}$ for $i = 1, \dots, n$ and $j = 1, \dots, m$, and values q and Q . (The value $p_{i,j} = \text{Prob}(X_i = a_{i,j})$.)*

Question: Is $\text{Prob}(\sum_{i=1}^n X_i \leq Q) \geq q$?

Theorem 4.1. *The problem of evaluating the cdf of convolutions of discrete variables is #P-hard even in the case that $m = 2$ and $p_{i,j} = \frac{1}{2}$ for all $i = 1, \dots, n$ and $j = 1, 2$.*

Proof. Let $A = a_1, \dots, a_n$, K , and B be the instance of the K th largest subset problem. For each $i = 1, \dots, n$, let X_i be the random variable that equals a_i with probability $\frac{1}{2}$ and equals 0 with probability $\frac{1}{2}$. Let $Q = B$, and let $q = \frac{K}{2^n}$. Then there are K or more distinct subsets of A with sum at most B if and only if $\text{Prob}(\sum_{i=1}^n X_i \leq Q) \geq q$. \square

We next establish the #P-hardness of the single-item stochastic inventory problem.

Theorem 4.2. *The single-item stochastic inventory control problem is #P-hard even in the case that all costs are linear.*

Proof. Let X_1, \dots, X_n , a_1, \dots, a_n , for $i = 1, \dots, n$ and $j = 1, 2 = m$, and values q and Q be the input for the Problem of evaluating the cdf of convolutions of discrete variables, where X_i takes value a_i with probability $\frac{1}{2}$ and takes value 0 with probability $\frac{1}{2}$. This problem is #P-hard as established in the proof of Theorem 4.1. We create an instance of the single-item stochastic inventory control problem as follows. Let $A = \sum_{i=1}^n a_i$. The Demand in period i is $D_i = X_i$. The cost of production in period 1 is $c_1(x) = (1 - q)x$. The cost of production in every other period $i \neq 1$ is $c_i(x) = Ax$. (This gives fractional costs. One can scale by 2^n to get integral costs.) The holding costs are 0 in each period, and the backorder costs are 0 in each period except for period n , in which $h_n(x) = -qx$ for $x < 0$. We now claim that there is a feasible solution for the Problem of Evaluating the cdf if and only if there is a feasible solution to the inventory control problem with expected cost at most $A - qQ$. To see that, we first note that any optimal policy will consist of ordering only in the first period. Any unsatisfied demand in later periods will be handled through backordering. This instance in turn is equivalent to newsvendor problem in which the cost c^m of ordering one unit too much is $(1 - q)$, and the cost of ordering one item too few is $c^\ell = q$. Therefore, the optimal decision is to produce the minimum amount S such that $\text{Prob}(D \leq S) \geq \frac{c^\ell}{c^\ell + c^m} = \frac{q}{1 - q + q} = q$, where D is a random variable describing the total demand (see, e.g., [SCB05], Section 8.2.1). The optimal order level is thus the minimum value x^* such that $\text{Prob}(\sum_{i=1}^n X_i \leq x^*) \geq q$. But calculating x^* is polynomially equivalent to evaluating the original cdf problem. \square

Remark. If costs are linear, then an optimal policy is “an order up to policy” in which there are values b_1, \dots, b_n , and the optimal policy is to order up to b_j in period j if inventory falls below period j . Now suppose that b'_1, \dots, b'_n is a proposed policy. (In the lingo of complexity analysis, b' is a certificate.) Note that already to *evaluate* the expected cost of this policy is #P-hard.

We note that many #P-hard enumeration problems have a fully polynomial randomized approximation scheme (FPRAS) including the following: counting Hamiltonian cycles in dense graphs, [DFJ98], counting knapsack solutions, [Dye03], counting Eulerian orientations of a directed graph, [MW95], counting perfect matchings in a bipartite graph, [JS89], and computing the permanent, [JSV04]. To the best of our knowledge, the only deterministic FPTAS for a #P-hard problem known up-to-date and published in the literature is the recent FPTAS of [Wei06] for counting independent sets in sparse graphs. (Counting independent sets of graphs of maximum degree 4 is known to be #P-complete, Theorem 3.1.5 in [Rot96].)

5 K -approximation sets and functions

A K -approximation algorithm for a minimization problem guarantees its output to be no more than K times the optimal solution. In this section we define K -approximation functions and K -approximation sets.

Definition 5.1. Let $K \geq 1$ and let $f : D \rightarrow \mathbb{R}_+$ be a function. We say that $\tilde{f} : D \rightarrow \mathbb{R}$ is a K -approximation function of f (K -approximation of f , in short) if for all $x \in D$ we have $f(x) \leq \tilde{f}(x) \leq Kf(x)$.

The following proposition follows directly from the definition of K -approximation functions.

Proposition 5.2. Let $K > 1$, let $f_1, f_2 : D \rightarrow \mathbb{R}_+$ be functions over domain D , let $\tilde{f}_1, \tilde{f}_2 : D \rightarrow \mathbb{R}_+$ be K -approximations of f_1, f_2 , respectively, let $g : D \rightarrow D$, and let $\alpha, \beta \in \mathbb{R}_+$. The following properties hold:

1. $\alpha + \beta\tilde{f}_1$ is a K -approximation of $\alpha + \beta f_1$,
2. $\tilde{f}_1 + \tilde{f}_2$ is a K -approximation of $f_1 + f_2$,
3. $\tilde{f}_1(g)$ is a K -approximation of $f_1(g)$,
4. $\tilde{f}_3(y) := \min_{x \in D} \{\tilde{f}_1(x) + \tilde{f}_2(x+y)\}$ is a K -approximation of $f_3(y) := \min_{x \in D} \{f_1(x) + f_2(x+y)\}$.

In order to get a polynomial time approximation scheme, we consider only a subset of all possible optimal inventory values, whose cardinality is polynomially bounded by the input size. Of course this can only be done by sacrificing accuracy in the final solution. We use the following definition.

Definition 5.3. Let $K > 1$ and let $f : [L, U] \rightarrow \mathbb{Z}_+$ be a monotone function. A K -approximation set of f is an ordered set $S = \{i_1 < \dots < i_r\}$ of integers satisfying the following two properties:

1. $L, U \in S \subseteq \{L, \dots, U\}$;
2. for each $k = 1$ to $r - 1$, if $i_{k+1} > i_k + 1$, then $\frac{f(i_k)}{K} \leq f(i_{k+1}) \leq Kf(i_k)$;

The canonical K -approximation set for a monotonically nondecreasing function is defined as follows:

$$i_1 \leftarrow L;$$

$$\text{for } k \geq 1, \text{ if } i_k < U, \text{ then } i_{k+1} \leftarrow \max\{i_k + 1, \max\{x \mid f(x) \leq Kf(i_k) \text{ and } x \leq U\}\}$$

The canonical K -approximation set for a monotonically nonincreasing function is defined analogously except that

$$i_{k+1} \leftarrow \max\{i_k + 1, \max\{x \mid f(x) \leq f(i_k)/K \text{ and } x \leq U\}\}.$$

One can determine each element of the canonical K -approximation set in $O(\log(U - L))$ time by using binary search. We have just proved the following lemma.

Lemma 5.4. Let $f : [L, U] \rightarrow \mathbb{Z}_+$ be a monotone function. For every $K > 1$ there exists a K -approximation set S of f of cardinality $O(\log_K f^{\max})$. Furthermore, it takes $O((1 + t_f) \log_K f^{\max} \log(U - L))$ time to construct this set, where t_f is the time needed to evaluate f .

We use K -approximation sets to construct approximations functions in the following way.

Definition 5.5. Let $K > 1$ and let $f : [L, U] \rightarrow \mathbb{Z}_+$ be a monotone function. Let S be a K -approximation set of f . A function \hat{f} defined as follows is called the approximation of f corresponding to S . For any integer $L \leq x \leq U$ and successive elements $i_k, i_{k+1} \in S$ with $i_k < x \leq i_{k+1}$ let

$$\hat{f}(x) := \begin{cases} f(x) & \text{if } x \in S; \\ \max\{f(i_k), f(i_{k+1})\} & \text{otherwise.} \end{cases}$$

Note that if we calculate the values of f on S in advance and store them in a sorted array $(x, f(x))$, then any query for the value of $\hat{f}(x)$, for any x , can be calculated in $O(\log |S|) = O(\log \log_K f^{\max})$ time. This is done by performing binary search over S to find the consecutive elements $i_k, i_{k+1} \in S$ such that $i_k < x \leq i_{k+1}$. The following proposition follows directly from the above definitions.

Proposition 5.6. *Let $K > 1$, let $f : [L, U] \rightarrow \mathbb{Z}_+$ be a nondecreasing function, and let S be a K -approximation set of f . If \hat{f} is the approximation of f corresponding to S , then \hat{f} is a nonnegative nondecreasing integer-valued step K -approximation function of f .*

We extend the definition of K -approximation sets to nonnegative unimodal functions in the following way.

Definition 5.7. *Let $K \geq 1$ and let $f : [L, U] \rightarrow \mathbb{Z}_+$ be a unimodal function with minimum value occurring at J . A K -approximation set of f if it can be expressed as the union $S = S_1 \cup S_2$ of a K -approximation set S_1 of the nonincreasing function $f : [L, J]$, and a K -approximation set S_2 of the nondecreasing function $f : [J, U]$. A function \hat{f} is called the approximation of f corresponding to S if it merges the approximations of f corresponding to S_1 and S_2 .*

Definition 5.5 holds for convex functions of one variables defined on an integer domain. Suppose now that f attains its minimum at x^* (if x^* is not unique, we set $x^* = \min\{\text{argmin}_{x \in [L, U]} f(x)\}$ to be the smallest such minimizer). Since f is convex, finding x^* can be done efficiently by standard binary search. Therefore, Lemma 5.4 and Proposition 5.6 hold for such functions as well. We summarize the results of this section as follows.

Theorem 5.8. *Let $f : [L, U] \rightarrow \mathbb{Z}_+$ be a unimodal function with a given $x^* = \text{arg min } f$. For every $K > 1$ there exists a K -approximation set S for f of cardinality $O(\log_K f^{\max})$ that can be constructed in $O(t_f \log_K f^{\max} \log(U - L))$ time, where t_f is the time needed to evaluate f . Furthermore, the approximation \hat{f} of f corresponding to S is a nonnegative integer-valued unimodal step K -approximation function of f , whose query time is $O(\log \log_K f^{\max})$.*

We denote an algorithm that calculates a canonical K -approximation set for a unimodal f with a given $x^* = \text{arg min } f$ by $\text{ApxSet}(K, f, x^*)$.

6 An FPTAS

In this section we develop an FPTAS for the inventory control problem defined in Section 3.

6.1 Preliminaries

Let $g_t(I_t)$ denote the optimal total expected cost for periods t, \dots, T , starting in period t with an inventory of I_t . Therefore, our goal is to calculate $z^* = g_1(0)$. Let $r_t(\bar{I}_t)$ be the expected cost for periods t, \dots, T , if the inventory at the *end* of period t is \bar{I}_t . It follows that for $t = 1, \dots, T$ we have

$$r_t(\bar{I}_t) = h_t(\bar{I}_t) + g_{t+1}(\bar{I}_t), \quad (2)$$

and

$$g_t(I_t) = \min_{x_t \in \mathbb{Z}_+} \{c_t(x_t) + \sum_{j=1}^{n_t} p_{t,j} r_t(I_t + x_t - d_{t,j})\}, \quad (3)$$

where $g_{T+1}(y) = 0$ for any y . Let us define an auxiliary function y_t as follows.

$$y_t(z) = \sum_{j=1}^{n_t} p_{t,j} r_t(z - d_{t,j}),$$

so $y_t(z)$ is the expected total cost for periods t, \dots, T after a procurement decision has been made in period t , and the procurement plus the inventory level is z . We note that

$$g_t(I_t) = \min_{x_t \in \mathbb{Z}_+} \{c_t(x_t) + y_t(I_t + x_t)\}. \quad (4)$$

Let us also define

$$\begin{aligned} R_t &= M_{t+1}r_t, \\ Y_t &= M_t y_t, \\ G_t &= M_t g_t. \end{aligned}$$

We state several basic properties of functions r_t, y_t, g_t, R_t, Y_t and G_t .

Proposition 6.1. *For every $t = 1, \dots, T$, functions r_t, g_t and y_t are convex over \mathbb{Z} .*

Proof. We prove this by backward induction. We consider first the base case of $t = T$. Since $g_{T+1} = 0$ we get that $r_T = h_T$, which is convex by assumption. Since the convexity of $r_T(z)$ induces the convexity of $r_T(z+d)$ for every constant d , and since a convex combination of convex functions is convex, we get that y_T is convex. We next prove that g_T is convex. From the definition of convex functions over discrete domains (see Section 2), it suffices to show that $2g(I) \leq g(I+1) + g(I-1)$ for all I . (For convenience we drop the subscript T .) Consider now a fixed value of I . Choose x' and x'' so that

$$\begin{aligned} g(I-1) &= c(x') + y(x' + I - 1); \\ g(I+1) &= c(x'') + y(x'' + I + 1). \end{aligned}$$

Case 1. $x' = x''$. Then $2g(I) \leq 2c(x') + 2y(x' + I) \leq 2c(x') + y(x' + I - 1) + y(x' + I + 1) = g(I-1) + g(I+1)$.

Case 2. $x' \geq x'' + 1$. Then $2g(I) \leq c(x' - 1) + y((x' - 1) + I) + c(x'' + 1) + y(x'' + 1 + I)$. Therefore, $2g(I) - g(I+1) - g(I-1) \leq c(x' - 1) + c(x'' + 1) - c(x') - c(x'')$, which is true because $x' > x''$ and c is convex.

Case 3. $x' \leq x'' - 1$. Then $2g(I) \leq c(x' + 1) + y((x' + 1) + I) + c(x'' - 1) + y(x'' - 1 + I)$. Therefore, $2g(I) - g(I+1) - g(I-1) \leq c(x' + 1) + c(x'' - 1) - c(x') - c(x'') + y((x' + 1) + I) + y(x'' - 1 + I) - y(x' + I - 1) - y(x'' + I + 1) \leq 0$. The latter inequality holds by the convexity of c and y and the fact that $x' \leq x'' - 1$.

Thus g_T is convex over \mathbb{Z} . Let us assume by induction that r_{t+1}, g_{t+1} and y_{t+1} are convex over \mathbb{Z} . We need to prove that r_t, g_t and y_t are convex over \mathbb{Z} as well. By assumption h_t is convex over \mathbb{Z} . By the induction hypothesis also is g_{t+1} . Hence, r_T is convex over \mathbb{Z} as the summation of two such functions. The rest of the proof for this case is similar to the base case. \square

We next show that all the values of $g_t(\cdot), r_t(\cdot)$ and $y_t(\cdot)$ over all inventory levels and time periods are rational numbers, and we bound the least common multiple of their denominators.

Proposition 6.2. *For every $t = 1, \dots, T$, functions $M_{t+1}r_t, M_t g_t, M_t y_t, R_t, Y_t$ and G_t are nonnegative, integer valued and convex over \mathbb{Z} .*

Proof. The proof for $M_{t+1}r_t$ and $M_t g_t$ is by induction. Since $r_T \equiv h_T$, and $g_{T+1} \equiv 0$, r_T is an integer function. Considering g_T we have

$$g_T(I_T) = \frac{1}{Q_T} \min_{x_T} \{ Q_T c_T(x_T) + \sum_{j=1}^{n_T} q_{T,j} r_T(I_T + x_T - d_{T,j}) \}.$$

Since c_T and r_T are integer-valued functions, $Q_T g_T$ is an integer-valued function as well. Assuming by induction that the statement holds for $t = k + 1$, we get immediately from (2) that the statement holds for r_k as well. From (3) we have

$$g_k(I_k) = \frac{1}{Q_k} \min_{x_k} \{ Q_k c_k(x_k) + \sum_{j=1}^{n_k} q_{k,j} r_k(I_k + x_k - d_{k,j}) \}.$$

Since c_k and $M_{k+1}r_k$ are both integer-valued functions, by the induction hypothesis, $M_k g_k$ is an integer-valued function as well. We conclude the proof for $M_{t+1}r_t$ and $M_t g_t$ by applying Proposition 6.1.

The proof for $M_t y_t$ is due to the fact that y_t is a convex combination of several r_t 's, and that Q_t is a common denominator of all the probabilities in time period t . The integrality and convexity of R_t, Y_t and G_t immediately follows. \square

The inventory I_t at the beginning of time period t following an optimal policy satisfies

$$-D^* \leq -\sum_{j=1}^{t-1} d_{j,n_j} \leq I_t \leq D^* - \sum_{j=1}^{t-1} d_{j,1} \leq D^*$$

for every $t = 1, \dots, T$. (To see this, note that the lower bound holds for any policy since $\sum_{j=1}^{t-1} (x_j - D_j) \geq -\sum_{j=1}^{t-1} d_{j,n_j}$. The upper bound holds for any optimal policy, since any such policy will order at most a total of D^* over the time horizon. Moreover, for every time period t , the inventory $I_t + x_t$ after the procurement decision has been made satisfies $-D^* \leq -\sum_{j=1}^{t-1} d_{j,n_j} \leq I_t + x_t \leq D^* - x_t - \sum_{j=1}^{t-1} d_{j,1} \leq D^*$.) Hence, x_t can be restricted to take values between 0 and D^* . The running time for computing the values of g_t and r_t by dynamic programming for all possible optimal inventory levels and for every period $t = 1, \dots, T$, is therefore $O(n^* T D^{*2})$, i.e., pseudo polynomial in the input size. By the discussion above we restrict without loss of generality the domain of r_t and g_t to be $[-D^*, \dots, D^*]$ and let $L = -D^*$, $U = D^*$.

We conclude this section by giving several properties of approximation functions.

Proposition 6.3. *Let $K > 1$ and f be an integer-valued function. If f' is a (general) K -approximation of f then $\lfloor f' \rfloor$ is an integer-valued K -approximation of f .*

This proposition is due to $f \leq \lfloor f' \rfloor \leq f' \leq Kf$, where the first inequality derives from the integrality of f and since f' is a K -approximation of f .

Proposition 6.4. *Let $K > 1$ and f be an integer-valued function. If \check{f} is a convex K -approximation of f with $\arg \min f = x^*$, then $\lfloor \check{f} \rfloor$ is a unimodal integer-valued K -approximation of f with the same $\arg \min$.*

Proof. The monotonicity of the floor function coupled with the convexity of \check{f} implies that $\lfloor \check{f} \rfloor$ is a unimodal function that is minimized at x^* . Function $\lfloor \check{f} \rfloor$ is an integer-valued K -approximation of f due to the previous proposition. \square

Proposition 6.5. *Let $K > 1$, f be a convex function, S be a K -approximation set of f , and \hat{f} be the approximation of f corresponding to S . Then the convex extension of \hat{f} induced by S is a convex K -approximation of f .*

Proof. This proposition is true because the convex extension of \hat{f} induced by S is the greatest convex function which does not lie above \hat{f} , and the fact that f itself is a convex function (i.e., the convex extension of \hat{f} induced by S is “sandwiched” between \hat{f} from above, and f from below). \square

6.2 Algorithm

In this section we give a formal description of our approximation scheme for r_t, g_t , and y_t for $t = 1, \dots, T$. The outline of the proof of correctness of the algorithm goes as follows. Since $g_{T+1} \equiv 0$ and h_T is given explicitly, we can calculate $r_T \equiv h_T$ by a single query to h_T and calculate y_T by performing n_T queries to h_T . Since $c_T(x)$ and $y_T(z+x)$ are convex functions in x for any fixed z , calculating $g_T(z)$ is done by performing binary search in the action space $[-D^* - z, D^* - z]$ to find an action x that minimizes $c_T(x) + y_T(z+x)$. This results in $\log U$ queries to c_T and y_T . As for calculating $g_t(z)$, in order not to perform overall $O(\log^{T+1-t} U)$ queries, which is exponential in the input size, our algorithm will efficiently compute a compressed approximation for g_t , so a query to it will cost only $O(\log \log_K \bar{U})$ time, for any

```

1: Function FPTAS( $\epsilon$ )
2: Let  $K := 1 + \frac{\epsilon}{2T}$  and  $\check{G}_{T+1} := 0$ 
3: for  $t := T$  downto 1 do
4:   Let  $\check{R}_t := M_{t+1}h_t + \check{G}_{t+1}$ 
5:   Let  $\check{Y}_t := Q_t \sum_{j=1}^{n_t} p_{t,j} \check{R}_t(z - d_{t,j})$ 
6:   Let  $G'_t(z) := \min_{x \in [-D^* - z, D^* - z]} \{M_t c_t(x) + \check{Y}_t(z + x)\}$ 
7:   Let  $x^* := \arg \min G'_t$  and  $\bar{G}_t := \lfloor G'_t \rfloor$ 
8:   Let  $S_t := \mathbf{ApxSet}(K, \bar{G}_t, x^*)$  and  $\hat{G}_t$  be its corresponding  $K$ -approximation function
9:   Let  $\check{G}_t$  be the convex extension of  $\hat{G}_t$  induced by  $S_t$ 
10: end for
11: Return  $\frac{\check{G}_1(0)}{M_1}$ 

```

Algorithm 1: An FPTAS.

t . The algorithm proceeds backwards from $t = T$ down to $t = 1$. In the $(T - t + 1)$ th iteration, given a previously calculated approximation for g_{t+1} , the algorithm approximates r_t , y_t , and g_t . The relative error of these later approximations is only slightly worse than the one of the given g_{t+1} , and the accumulated error throughout the execution of the algorithm is under control.

6.3 Analysis

Let $\epsilon > 0$ be given, where we seek a $(1 + \epsilon)$ -approximation, as in (1). In this section we show that Algorithm 1 produces the required approximation, and runs in time that is polynomial in both the input size and $\frac{1}{\epsilon}$.

Lemma 6.6. *Algorithm 1 computes \check{G}_t , which is a convex K^{T+1-t} -approximation of G_t , for every $t = 1, \dots, T + 1$.*

Proof. We first note that \bar{G}_t is an integer-valued unimodal function. To see this, note that by Proposition 6.1, G'_t is a convex function, and the floor function is monotone.

We prove the approximation ratio by backwards induction. We consider first the base case of time period T . Since $G_{T+1} = \check{G}_{T+1} = 0$, we get from the definition of \check{R} that $\check{R}_T = R_T$. Hence, $\check{Y}_T = Y_T$ and $G'_T = G_T$. Therefore, by Proposition 6.2, G'_T is an integer-valued convex function, and $\bar{G}_T = G_T$. By Theorem 5.8, \hat{G}_T is an integer-valued K -approximation function of G_T . By Proposition 6.5, \check{G}_T is a (not necessarily integer-valued) convex K -approximation function of G_T as needed.

We assume inductively that \check{G}_{t+1} is a convex K^{T-t} -approximation of G_{t+1} . By the second property of Proposition 5.2 we get that \check{R}_t is a (convex) K^{T-t} -approximation of R_t ($M_{t+1}h_t$ is a K^{T-t} -approximation of itself). By the first two properties of the same proposition, \check{Y}_t is a convex K^{T-t} -approximation of Y_t . By the last three properties of the same proposition and Proposition 6.1, we get that G'_t is a convex K^{T-t} -approximation function of G_t . Due to Proposition 6.4, \bar{G}_t is a convex K^{T-t} -approximation function of G_t , i.e.,

$$G_t \leq \bar{G}_t \leq K^{T-t} G_t. \quad (5)$$

By Theorem 5.8, \hat{G}_t is a K -approximation of \bar{G}_t . Therefore, by (5) we have

$$G_t \leq \bar{G}_t \leq \hat{G}_t \leq K \bar{G}_t \leq K^{T+1-t} G_t;$$

so \hat{G}_t is a K^{T+1-t} -approximation of G_t . We conclude the proof by applying Proposition 6.5. \square

We conclude this section by proving that Algorithm 1 is an FPTAS for our problem.

Theorem 6.7. *Algorithm 1 gives an FPTAS for the single-item discrete stochastic inventory control problem when $K = 1 + \frac{\epsilon}{2T}$, for any $\epsilon < 1$.*

Proof. By Lemma 6.6, the approximation $\hat{g}_1(0) := \frac{\check{G}_1(0)}{M_1}$ of the optimal total expected cost in periods $1, \dots, T$ starting in period 1 with an inventory of 0 satisfies $g_1(0) \leq \hat{g}_1(0) \leq K^T g_1(0)$. Setting $K = 1 + \frac{\epsilon}{2T}$ gives $g_1(0) \leq \hat{g}_1(0) \leq (1 + \frac{\epsilon}{2T})^T g_1(0)$. From the inequality $(1 + \frac{x}{n})^n \leq 1 + 2x$, which holds for every $0 \leq x \leq 1$, and since the optimal solution is $g_1(0)$, we get that $z^* \leq \hat{g}_1(0) \leq (1 + \epsilon)z^*$.

We now compute upper bounds on the values of the functions computed during the execution of the algorithm. Let

$$A = 2T \max_t \{c_t(D^*), h_t(D^*), h_t(-D^*)\} \geq \sum_{t=1}^T (c_t(D^*) + \max\{h_t(D^*), h_t(-D^*)\})$$

be an upper bound for the values of the r_t 's and g_t 's. Let

$$B = K^T M_1 A.$$

Due to Lemma 6.6, B is an upper bound for the values of the various $\check{R}_t, \check{Y}_t, G'_t, \bar{G}_t$'s and \check{G}_t 's. Therefore, B serves as an upper bound for the functions considered throughout the algorithm. Note that Assumption 3.4 implies that $\log B$ is polynomially bounded by the input size. Since $U = O(D^*)$, $\log U$ is also polynomially bounded by the input size.

We next consider the running time of our algorithm. It consists of T iterations. We next analyze the running time for iteration t only. Step 4 is executed in constant time. A query to \check{R}_t is done by performing a single query to each of h_t and \check{G}_t . The former takes t_h time, and the later takes $O(\log \log_K B)$ time due to Theorem 5.8. Step 5 is done in constant time, and each query to \check{Y}_t is performed in $O(n_t(t_h + \log \log_K B))$ time. Step 6 is done in constant time. A calculation of $G'_t(z)$ is carried out by performing binary search over $[-D^* - z, D^* - z]$ since $M_t c_t(x) + \check{Y}_t(z+x)$ is a convex function of x for every fixed z . Hence, a query of $G'_t(\cdot)$ is done in $O(\log U(t_c + n_t(t_h + \log \log_K B)))$ time. As for Step 7, due to Proposition 6.1 G'_t is convex. So finding x^* is done by binary search over $[-D^*, D^*]$, resulting in a total time of $O(\log^2 U(t_c + n_t(t_h + \log \log_K B)))$. By Theorem 5.8, Step 8 is performed in $O(\log^2 U \log_K B(t_c + n_t(t_h + \log \log_K B)))$. Since the elements in S are kept sorted, the construction of the convex extension in Step 9 is done in linear time in $|S|$ [PS85]. Therefore, the running time of each iteration of the for loop is dominated by Step 8. The the total running time of the algorithm for iteration t is

$$O(T \log^2 U \log_K B(t_c + n_t(t_h + \log \log_K B))).$$

Note that $O(\log_K B) = O(T + \log_K(M_1 A))$. Without loss of generality, we can assume that $\epsilon < 1$ and thus that $K < 2$. So $O(\log_K Y) = O(\frac{1}{K-1} \log Y)$ for every Y . Replacing K with $1 + \frac{\epsilon}{2T}$ we obtain $O(\log_K Y) = O(\frac{T}{\epsilon} \log Y)$. In this way we get that $O(\log_K B) = O(\frac{T}{\epsilon} \log(M_1 A))$.

Replacing a query time of c_t and h_t by t_f we conclude that the running time of the algorithm is

$$O\left(\frac{T^2 \log^2 U \log(M_1 A)}{\epsilon} [n^*(t_f + \log(\frac{T \log(M_1 A)}{\epsilon}))]\right),$$

which is polynomial in the binary size of the input data and $\frac{1}{\epsilon}$. The dependance on T is almost quadratic, and that of $\frac{1}{\epsilon}$ is almost linear. \square

7 Extensions

7.1 Capacitated version

Convex cost functions can model capacity limits by making the procurement cost sufficiently high beyond these limits, while preserving convexity. Thus, our results hold also for single-item stochastic capacitated inventory control problems with discrete demands.

7.2 The lost sales model

So far, we have assumed that excess demand is backlogged. We now consider the case when the excess demand is lost. Even if the costs are convex, the value function g_t is not necessarily convex. However, we can transform the lost sales case to the previous convex cost case if all costs are linear, and under a reasonable assumption on the cost of lost sales, to be specified next.

We consider a problem with lost sales, in which $I_t \geq 0$ in each period t , and where the lost sales in period t is $L_t = \max(0, -\bar{I}_t)$. In addition, the linear costs of production, inventory, and lost sales in period t are $c_t x_t$, $h_t \bar{I}_t$, and $\ell_t L_t$ respectively. We also assume that $\ell_t \geq c_{t+1} - h_t$ for $t = 1, \dots, T$. Note that this condition is typically true in practice since typically $\ell_t \geq c_{t+1}$ in practice. [Zip00], pages 386-387 showed under this assumption that the optimal value function is convex, and a base stock policy is optimal. We next give an alternative proof by transforming this lost sales instance to an instance of our original stochastic inventory control problem with convex costs and backordering. We transform the lost sales model by splitting period t in the original problem into two periods in the transformed problem, denoted as periods $2t - 1$ and $2t$. Period $2t - 1$ corresponds to period t of the original problem. Production in period $2t$ in the transformed problem corresponds to lost sales in period t of the original problem. As such we define costs $c'_t(x)$ and $h'_t(\bar{I}_t)$ for the transformed problem as follows: $c'_{2t-1}(x) = c_t x$, $h'_{2t-1}(I) = 0$, $D'_{2t-1} = D_t$; $c'_{2t}(x) = \ell_t x$,

$$h'_{2t}(I) = \begin{cases} h_t I & \text{for } I \geq 0; \\ -MI & \text{otherwise,} \end{cases}$$

and $D'_{2t} = 0$, where M is a very large number polynomially bounded by the input size (e.g., $M = D^* \max_t c_t$). Note that h'_{2t} is convex.

Suppose that we have a feasible policy to the original problem. For any realization of the demand and for any solution with cost C , we can transform a feasible solution for the original problem into a feasible solution for the transformed problem with the same cost by setting $x'_{2t} = x_t$ and $x'_{2t+1} = L_t$.

We may restrict attention to solutions for the transformed problem that have $x'_{2t} =$ backordering in period $2t - 1$. There is no additional production in period $2t$ since any additional unit would need to be inventoried to period $2t + 1$, and due to our assumptions above the cost would be no higher by producing the unit in period $2t + 1$.

Any solution satisfying the condition stated in the above paragraph leads to a solution for the original problem with the same cost.

7.3 Discounted version

In the discounted single-item stochastic inventory control problem, we are also given a rational discounting factor $0 < \lambda = \frac{\lambda_1}{\lambda_2} < 1$, where λ_1 and λ_2 are positive integers. In this version the objective function changes to

$$z^* = \min E_D \left(\sum_{t=1}^T \lambda^{t-1} (c_t(x_t) + h_t(I_t + x_t - D_t)) \right).$$

It is easy to adapt our algorithm to this case. The recursive formula (2) for r_t changes to

$$r_t(\bar{I}_t) = h_t(\bar{I}_t) + \lambda g_{t+1}(\bar{I}_t), \tag{6}$$

while (3) remains intact. We need to account for the fact that $r_t(\bar{I}_t)$ is not necessarily integer valued even if $g_{t+1}(\bar{I}_t)$ is integer valued. However, by multiplying (6) by λ_2 , i.e., by changing Q_t to be $Q_t = \lambda_2 \sum_{j=1}^{n_t} q_{t,j}$, we recover integrality, and we can use the same algorithm. The analysis of the running time can be easily extended. The only change is that M_1 is multiplied by λ_2^T , and therefore the running time of the entire algorithm is multiplied by $O(T \log \lambda_2)$. Therefore, it remains polynomial in the binary size of the data.

7.4 Disposal at a cost

We can also approximate a version of the problem where the procurement cost function is defined over \mathbb{Z}_- as well. If the procurement level is positive, then the actual procurement is occurred, otherwise, a disposal cost is charged. We replace Assumption 3.2 by the following assumption.

Assumption 7.1. *The procurement cost function c_t is nonnegative convex over \mathbb{Z} and $c_t(0) = 0$ for every $t = 1, \dots, T$.*

By extending the value of the procurement level of an optimal policy to be in the interval $[-D^*, D^*]$ we get that our results carry over for this case as well.

7.5 Constant lead time

Under general lead times, the value function is multivariate. It is well known that this dynamic program can be transformed into a single variate dynamic program (the state corresponds to inventory position, which is defined as the inventory on-hand and all outstanding inventory) of the same form as the one presented in Section 3. It is easy to show that this transformation preserves the approximation ratio and as a result it suffices to find an FPTAS for this single variate dynamic program. If $L > 0$ is an arbitrary lead time, then the underlying demand distribution of the transformed problem is $\bar{D}_t = \sum_{\hat{t}=t}^{t+L-1} D_{\hat{t}}$. The presented FPTAS requires that we know $\text{Prob}[\bar{D}_t = \bar{d}_{t,i}]$, which is a convolution of L distributions. As a result, computing these probabilities takes $(n^*)^L$ time. If L is 2 or 3 (or any other constant value), then the term $(n^*)^L$ is polynomial, and the algorithm is an FPTAS. If L is not constrained to be small (e.g., $L = T/4$), then the running time is exponentially large. In the latter case, our algorithm is not an FPTAS. It is an open question whether one can modify the approach and create an FPTAS for the problem in which the lead times are permitted to be a fraction of T .

Note that due to Theorem 4.1, in the presence of lead times it is hard to compute even a myopic policy that aims to minimize the period cost a lead time ahead.

7.6 Non-exact evaluation of cost functions

In Assumption 3.4, we require that c and h be evaluated in polynomial time. We can weaken this assumption as follows.

Assumption 7.2. *For every $\delta \geq 0$ and for every time period t there exist convex integer-valued convex functions \tilde{c}_t^δ and \tilde{h}_t^δ such that*

$$\frac{|\tilde{c}_t^\delta(x) - c_t(x)|}{c_t(x)} \leq \delta, \quad \frac{|\tilde{h}_t^\delta(x) - h_t(x)|}{h_t(x)} \leq \delta,$$

for every x , and these functions can be evaluated in polynomial time in the input size and $1/\delta$.

This assumption is equivalent to the statement that for every $K > 1$, c_t and h_t have a two-sided K -approximation.

Definition 7.3. *Let $K > 1$ and let $f : D \rightarrow \mathbb{R}$ be a function. We say that $\tilde{f} : D \rightarrow \mathbb{R}$ is a two-sided K -approximation of f if for all $x \in D$ we have $f(x)/K \leq \tilde{f}(x) \leq Kf(x)$.*

Let f be either c_t or h_t . By assumption 7.2, for every $\bar{K} > 1$ there exists an integer-valued convex function $\tilde{f}_{\bar{K}}$ such that $f(x)/\bar{K} \leq \tilde{f}_{\bar{K}}(x) \leq \bar{K}f(x)$ for every x . Let $t_{\tilde{f}}$ be the time needed to evaluate $\tilde{f}_{\bar{K}}$ on x . By Assumption 7.2, $t_{\tilde{f}}$ is polynomially bounded in the size of the problem and in $\frac{1}{\bar{K}-1}$. We set $\bar{K} = \sqrt[3]{K}$ so that for every pair x_1, x_2 of succeeding elements in the \bar{K} -approximation set S of $\tilde{f}_{\bar{K}}$ we get

$$\frac{f(x_2)}{f(x_1)} \leq \frac{\bar{K}\tilde{f}_{\bar{K}}(x_2)}{\frac{1}{\bar{K}}\tilde{f}_{\bar{K}}(x_1)} = \bar{K}^2 \frac{\tilde{f}_{\bar{K}}(x_2)}{\tilde{f}_{\bar{K}}(x_1)} \leq \bar{K}^3 = K.$$

Theorem 5.8 remains true with the exception that the approximation corresponding to S is a two-sided K -approximation. The FPTAS for the inventory control problem remains the same, its analysis is identical, and the resulting approximation functions are two-sided $(1 + \epsilon)$ -approximations. In particular $\frac{g_1(0)}{1+\epsilon} \leq \hat{g}_1(0) \leq (1 + \epsilon)g_1(0)$.

Note that $(1 + \epsilon)\hat{g}_1(0)$ serves as a one-sided $(1 + \epsilon)^2$ -approximation for $g_1(0)$. In order to get a one-sided $(1 + \epsilon)$ -approximation we change Step 2 of Algorithm 1 such that $K := 1 + \frac{\sqrt{1+\epsilon}-1}{2T}$, so the algorithm calculates \hat{g}_1 , which is a two-sided $(\sqrt{1 + \epsilon})$ -approximation of g_1 . Therefore, $(\sqrt{1 + \epsilon})\hat{g}_1$ is a one-sided $(1 + \epsilon)$ -approximation of g_1 as needed.

8 Conclusions and Future Research

We presented the first FPTAS for the single-item stochastic inventory control problem. Other recent developments in approximation algorithms for stochastic dynamic and multistage programs are based on gradients or sampling. Our framework is based on the notion of approximation sets and functions. We still use the standard optimality equation or recursion; however, we consider only polynomially many states. Our algorithm relies on the convexity of the value function.

We presented extensions to the basic model that do not require substantial modifications to the algorithm. It is an interesting open question of whether there is an FPTAS if one relaxes the assumption of constant lead time. Yet another interesting open extension is the consideration of the infinite time horizon problem under stationary data.

In [DGGJ03], the authors investigate classes of counting problems that are irreducible under approximation-preserving reductions. One of these classes is the class of counting problems that admit an FPRAS. It is therefore interesting in this context to investigate the class of counting problems that admit an FPTAS.

Acknowledgments. We thank the referees for their detailed comments which helped us to improve the presentation of this paper considerably. We thank Retsef Levi for pointing out that our original proof of #P-hardness for stochastic inventory control also implies the #P-hardness of evaluating cdfs for convolutions of discrete variables. Last, we thank Oded Goldreich for inspiring discussions and Leslie Ann Goldberg for pointing out references.

References

- [DFJ98] M. Dyer, A. Frieze, and M. Jerrum. Approximately counting Hamilton paths and cycles in dense graphs. *SIAM Journal on Computing*, 27:1262–1272, 1998.
- [DGGJ03] M. Dyer, L.A. Goldberg, C. Greenhill, and M. Jerrum. The relative complexity of approximate counting problems. *Algorithmica*, 38:471–500, 2003.
- [DGV08] B.C. Dean, M. X. Goemans, and J. Vondrák. Approximating the stochastic knapsack problem: the benefit of adaptivity. *Mathematics of Operations Research*, 33:945–964, 2008.
- [Dye03] M. Dyer. Approximate counting by dynamic programming. In *Proceedings of the 35th Annual ACM Symposium on the Theory of Computing*, pages 693–699, 2003.
- [FZ84a] A. Federgruen and P.H. Zipkin. Computational issues in an infinite-horizon, multi-echelon inventory model. *Operations Research*, 32:818–836, 1984.
- [FZ84b] A. Federgruen and P.H. Zipkin. An efficient algorithm for computing optimal (s,S) policies. *Operations Research*, 32:1268–1286, 1984.
- [GJ79] M.R. Garey and D.S. Johnson. *Computers and Intractability: A Guide to the Theory of NP-completeness*. W. H. Freeman, 1979.
- [JS89] M. Jerrum and A. Sinclair. Approximating the permanent. *SIAM Journal on Computing*, 18:1149–1178, 1989.

- [JSV04] M. Jerrum, A. Sinclair, and E. Vigoda. A polynomial-time approximation algorithm for the permanent of a matrix with nonnegative entries. *Journal of the ACM*, 51(4):671–697, 2004.
- [LPRS07] R. Levi, M. Pal, R. Roundy, and D.B. Shmoys. Approximation algorithms for stochastic inventory control models. *Mathematics of Operations Research*, 32:284–302, 2007.
- [LRS07] R. Levi, R. Roundy, and D.B. Shmoys. Provably near-optimal sampling-based policies for stochastic inventory control models. *Mathematics of Operations Research*, 32:821–838, 2007.
- [LRST08] R. Levi, R. Roundy, D.B. Shmoys, and V.A. Truong. Approximation algorithms for capacitated stochastic inventory control problems. *Operations Research*, 56:1186–1199, 2008.
- [MW95] M. Mihail and P. Winkler. On the number of Eulerian orientations of a graph. *Algorithmica*, 16:402–414, 1995.
- [Por02] E.L. Porteus. *Foundations of Stochastic Inventory Theory*. Stanford University Press, 2002.
- [PS85] F.P. Preparata and M.I. Shamos. *Computational Geometry: An Introduction*. Springer, 1985.
- [RM00] R. Roundy and J.A. Muckstadt. Heuristic computation of periodic-review base stock inventory policies. *Management Science*, 46:104–109, 2000.
- [Rot96] D. Roth. On the hardness of approximate reasoning. *Artificial Intelligence*, 82:273–302, 1996.
- [SCB05] D. Simchi-Levi, X. Chen, and J. Bramel. *The Logic of Logistics: Theory, Algorithms, and Applications for Logistics and Supply Chain Management*. Springer-Verlag, second edition, 2005.
- [SO95] H. Safer and J. Orlin. Fast approximation schemes for multi-criteria flow, knapsack, and scheduling problems. Working paper 3757-95, Sloan School of Business, Massachusetts Institute of Technology, 1995.
- [SS05] C. Swamy and D.B. Shmoys. Sampling-based approximation algorithms for multi-stage stochastic optimization. In *46th Annual IEEE Symposium on Foundations of Computer Science*, pages 357–366, Pittsburgh, PA, 2005.
- [SS06] D.B. Shmoys and C. Swamy. An approximation scheme for stochastic linear programming and its application to stochastic integer programs. *Journal of the ACM*, 53:973–1012, 2006.
- [VW01] C.P.M. Van Hoesel and A.P.M. Wagelmans. Fully polynomial approximation schemes for single-item capacitated economic lot-sizing problems. *Mathematics of Operations Research*, 26:339–357, 2001.
- [Wei06] D. Weitz. Counting independent sets up to the tree threshold. In *Proceedings of the 38th Annual ACM Symposium on Theory of Computing*, 2006.
- [Woe00] G.J. Woeginger. When does a dynamic programming formulation guarantee the existence of a fully polynomial time approximation scheme (FPTAS)? *INFORMS Journal on Computing*, 12:57–75, 2000.
- [ZF91] Y.Z. Zheng and A. Federgruen. Finding optimal (s,S) policies is about as simple as evaluating a single policy. *Operations Research*, 39:654–665, 1991.
- [Zip00] P.H. Zipkin. *Foundations of Inventory Management*. McGraw-Hill, 2000.