

Chance Constrained Knapsack Problem with Random Item Sizes

Vineet Goyal* R. Ravi†

June 9, 2009

Abstract

We consider a stochastic knapsack problem where each item has a known profit but a random size. The goal is to select a profit maximizing set of items such that the probability of the total size of selected items exceeding the knapsack size is at most a given threshold. This problem finds applications in capital investment problems where a set of projects needs to be selected from the available pool of projects to invest a given amount of capital. The investment requirement for each project is typically uncertain in the selection phase and is known only after it has been selected and started.

We present an FPTAS for the case when each item size is normally distributed and independent of other items. We present a parametric LP formulation and show that it is a good approximation of the chance-constrained stochastic knapsack problem. Furthermore, we give a polynomial time algorithm to round any fractional solution of the parametric LP to obtain an integral solution whose profit is within $(1 + \epsilon)$ -factor of the objective value of the fractional solution for any $\epsilon > 0$.

1 Introduction

We consider the following stochastic variant of the classical knapsack problem. We are given n items with profits p_1, p_2, \dots, p_n , a knapsack size B and a reliability level $0 \leq \rho \leq 1$. Item i has a random size S_i distributed according to a known distribution and independent of the sizes of other items. The goal is to select a subset \mathcal{S} of items such that,

$$Pr \left(\sum_{i \in \mathcal{S}} S_i \leq B \right) \geq \rho, \quad (1.1)$$

and the profit is maximized. We refer to (1.1) as the *chance-constraint* and the problem as the chance-constrained knapsack problem. Our model finds applications in problems where there is only one stage of decision making under uncertainty. For instance, consider the following capital investment problem where a central planner needs to select a set of projects to invest the available capital from a universe of projects where each project has an uncertain investment requirement and becomes known only after the project has been selected and started. The central planner would ideally like to invest in the set of projects that have a high profit or return while not exceeding the budget constraint. The solution might be highly conservative if we require that the budget is not exceeded in any possible realization of the investment requirements of selected projects. A chance-constrained model overcomes this drawback by allowing constraint violation for a small fraction of the unlikely realizations.

We present an FPTAS for the chance-constrained knapsack problem when item sizes are normally distributed and independent of other items. It is known [1] that in the case of normally distributed item sizes the chance-constraint can be formulated as a 0-1 conic program. However, we show that the integrality gap

*Operations Research Center, Massachusetts Institute of Technology, Cambridge, MA 02139. Email: goyalv@mit.edu. Supported in part by NSF grant CCF-0430751.

†Tepper School of Business, Carnegie Mellon University, Pittsburgh PA 15232. Email: ravi@andrew.cmu.edu. Supported in part by NSF grant CCF-0430751.

of the conic formulation is large. We reformulate the problem as a parametric LP and give a rounding algorithm that rounds any fractional solution to a $(1 + \epsilon)$ -approximate integral solution for any constant $\epsilon > 0$ in time polynomial in the input size and $\frac{1}{\epsilon}$.

Several stochastic variants of the classical knapsack problem have been studied in literature. Henig [5] and Carraway et al. [2] consider the stochastic variant where item sizes are known but the profit of each item is distributed normally and independent of others and the goal is to maximize the probability that the profit is at least a given threshold. The authors present dynamic programming and branch and bound heuristics to solve this problem to optimality. Papastavrou et al. [8] and Kleywegt et al. [7] consider a variant called as stochastic and dynamic knapsack problem where items arrive online according to some stochastic process - the size and profit of each item is known only after the item arrives and you are required to decide whether to select the item or not when it arrives. Dean, Goemans and Vondrak [3] study the benefit of adaptivity in the online stochastic knapsack problem and give a polynomial time non-adaptive policy that is within a factor 4 of the optimal adaptive policy. Note that an adaptive policy depends on the remaining knapsack capacity while a non-adaptive policy does not.

Kleinberg et al. [6] and Goel and Indyk [4] consider a chance-constrained stochastic knapsack problem similar to the one considered in this paper. Kleinberg et al. [6] consider the case where item sizes have a Bernoulli-type distribution (with only two possible sizes for each item), and provide an $O(\log 1/(1 - \rho))$ -approximation algorithm where ρ is the threshold probability. Goel and Indyk [4] provide a PTAS for the case when item sizes have Poisson or exponential distribution. However, the algorithm in [4] violates the chance constraint by a factor of $(1 + \epsilon)$. In contrast, we present an FPTAS for the case when item sizes are normally distributed while satisfying the chance-constraint strictly.

2 Conic Integer Formulation

We consider the case when each item j has a normally distributed size with mean a_j and standard deviation σ_j independent of the other items. Let x_j denote whether item j has been selected or not. Then the stochastic knapsack problem can be formulated as follows:

$$\max \left\{ \sum_{j=1}^n p_j x_j \mid \Pr \left(\sum_j S_j x_j \leq B \right) \geq \rho, x_j \in \{0, 1\}^n \right\}. \quad (2.1)$$

We simplify the probabilistic constraint as follows:

$$\Pr \left(\sum_j S_j x_j \leq B \right) = \Pr \left(\frac{\sum_j (S_j x_j - a_j x_j)}{\sqrt{\sum_j \sigma_j^2 x_j^2}} \leq \frac{B - \sum_j a_j x_j}{\sqrt{\sum_j \sigma_j^2 x_j^2}} \right). \quad (2.2)$$

Let,

$$Z = \left(\frac{\sum_j (S_j x_j - a_j x_j)}{\sqrt{\sum_j \sigma_j^2 x_j^2}} \right).$$

Since the item sizes are normally distributed and independent of other items, Z is a standard normal variable with mean 0 and standard deviation 1. Let ϕ denote the cumulative distribution function of the standard normal variate. Therefore, the probabilistic constraint can be rewritten as:

$$\Pr \left(Z \leq \frac{B - \sum_j a_j x_j}{\sqrt{\sum_j \sigma_j^2 x_j^2}} \right) \geq \rho \Rightarrow \frac{B - \sum_j a_j x_j}{\sqrt{\sum_j \sigma_j^2 x_j^2}} \geq \phi^{-1}(\rho),$$

where, $\phi^{-1}(\rho)$ is positive if $\rho > 0.5$. Therefore, we can reformulate (2.1) as follows:

$$\begin{aligned} & \max \sum_{j=1}^n p_j \cdot x_j \\ & \phi^{-1}(\rho) \sqrt{\sum_j \sigma_j^2 x_j^2} + \sum_j a_j x_j \leq B \\ & x_j \in \{0, 1\}, \forall j = 1, \dots, n \end{aligned} \quad (2.3)$$

If we relax the 0-1 constraints on x_j to $0 \leq x_j \leq 1$ for all $j = 1, \dots, n$, the formulation in (2.3) becomes a second order conic program and can be solved in polynomial time. However, the integrality gap of the conic relaxation is $\Omega(\sqrt{n})$.

Large Integrality Gap. Consider the following instance: $p_j = \sigma_j = 1, a_j = 1/\sqrt{n}$ for all $j = 1, \dots, n, B = 3, \rho = 0.95$. Any integral solution can at most two items; therefore, the integral profit is at most 2. Now, consider the fractional solution $x_j = \frac{1}{\sqrt{n}}$. Then,

$$\sum_{j=1}^n a_j x_j + \phi^{-1}(\rho) \sqrt{\sum_{j=1}^n \sigma_j^2 x_j^2} = 1 + \phi^{-1}(\rho) < 3.$$

Therefore, the fractional solution is feasible and the optimal fractional profit is at least \sqrt{n} which shows that the integrality gap of the conic formulation is $\Omega(\sqrt{n})$.

3 Parametric LP Reformulation

We reformulate the second order conic program as a parametric LP and obtain a fully polynomial time approximation scheme for the chance constrained knapsack problem. Suppose we know that the sum of mean sizes of the items selected in an optimal solution is μ^* . Then, the conic constraint in (2.3) can be expressed as,

$$\begin{aligned} \sum_j a_j x_j & \leq \mu^* \\ (\phi^{-1}(\rho))^2 \left(\sum_j \sigma_j^2 x_j^2 \right) & \leq (B - \mu^*)^2 \end{aligned} \quad (3.1)$$

Since $x_j^2 = x_j$ for $x_j \in \{0, 1\}$, we can simplify (3.1) as:

$$(\phi^{-1}(\rho))^2 \left(\sum_j \sigma_j^2 x_j \right) \leq (B - \mu^*)^2.$$

Therefore, we can formulate the chance constrained knapsack problem as the following parametric 2-dimensional knapsack problem where μ is the parameter corresponding to the total mean size of the selected items. We consider powers of $(1 + \epsilon)$, i.e., $(1 + \epsilon)^j, j = 0, \dots, \log_{(1+\epsilon)} B$ for some constant $\epsilon > 0$, as different choices of the parameter μ . Therefore, the number of different choices of μ is $O\left(\frac{\log B}{\epsilon}\right)$ which is polynomial in the input size.

We also guess the value of optimal profit OPT by considering powers of $(1 + \epsilon)$. Let $P = \sum_{j=1}^n p_j$; we consider $O\left(\frac{\log P}{\epsilon}\right)$ different choices of OPT. At most $\frac{1}{\epsilon}$ items can have profit greater than ϵOPT .

Therefore, for each guess of $\text{OPT} = (1 + \epsilon)^j$ we consider all subsets of size at most $\frac{1}{\epsilon}$ of the items that have size more than ϵOPT to include in the solution. For each guess O of OPT and each choice of subset of items of size more than ϵOPT , we solve a subproblem $\Pi(S_1, S_2, O)$ where S_2 is the set of items whose profit is more than $\epsilon \cdot O$ and are included in the final solution and we are required to choose a subset of items from $S_1 \subset [n]$ that maximizes the total profit. Let $\Pi(S_1, S_2, O, \mu)$ denote the problem where the total mean size of all items selected from S_1 is at most μ . Therefore, we can formulate $\Pi(S_1, S_2, O, \mu)$ as the following 2-dimensional knapsack problem:

$$\begin{aligned} \max \quad & \sum_{j \in S_1} p_j x_j + \sum_{j \in S_2} p_j \\ & \sum_{j \in S_1} a_j x_j \leq \mu \end{aligned} \quad (3.2)$$

$$\begin{aligned} (\phi^{-1}(p))^2 \left(\sum_{j \in S_1} \sigma_j^2 x_j \right) & \leq \left(B - \mu - \sum_{j \in S_2} \mu_j \right)^2 - (\phi^{-1}(p))^2 \left(\sum_{j \in S_2} \sigma_j^2 \right) \\ x_j & \in \{0, 1\} \end{aligned} \quad (3.3)$$

Algorithm \mathcal{A} for the Chance-constrained Knapsack Problem is described in Figure 1 and Algorithm $\mathcal{A}(\Pi)$ for $\Pi(S_1, S_2, O, \mu)$ is described in Figure 2.

Algorithm \mathcal{A} for Chance-constrained Knapsack Problem.

Input: Given n items where item j has profit p_j and a normally distributed size with mean a_j and standard deviation σ_j , knapsack size B , reliability level ρ and a constant $\epsilon > 0$. Let $p_m = \min_{j \in [n]} p_j$, $P = \sum_{j \in [n]} p_j$.

Initialize $N_1 = \lfloor \log_{1+\epsilon} p_m \rfloor$, $N_2 = \lceil \log_{1+\epsilon} P \rceil$, $x_{\mathcal{A}} = 0$, $P_{\mathcal{A}} \leftarrow 0$.

1. For $t = N_1, \dots, N_2$,
 - (a) Let $O = (1 + \epsilon)^t$ and let $S_{\epsilon} = \{j \in [n] \mid p_j \geq \epsilon \cdot O\}$.
 - (b) For each set $S \subset S_{\epsilon}$ such that $|S| < \frac{1}{\epsilon}$,
 - i. Solve $\Pi([n] \setminus S_{\epsilon}, S, O)$ and let x_S denote the integral solution returned by $\mathcal{A}(\Pi)$.
 - ii. If $P_{\mathcal{A}} < p^T x_S$, then

$$\begin{aligned} x_{\mathcal{A}} & \leftarrow x_S \\ P_{\mathcal{A}} & \leftarrow p^T x_S \end{aligned}$$

2. Return the solution $x_{\mathcal{A}}$.

Figure 1: Algorithm \mathcal{A} for Chance-constrained Knapsack Problem

In the following lemma, we show that we can find a good integral solution to the problem $\Pi(S_1, S_2, O, \mu)$.

Lemma 3.1 Consider the problem $\Pi(S_1, S_2, O, \mu)$ such that $p_j \leq \epsilon \cdot O$ for all $j \in S_1$. If P^* is the optimal profit for $\Pi(S_1, S_2, O, \mu)$, then there is a polynomial time algorithm to find a feasible set of items whose profit is at least $(P^* - 2\epsilon \cdot O)$.

Proof: Consider the 2-dimensional knapsack formulation of $\Pi(S_1, S_2, O, \mu)$ and consider the basic optimal solution \tilde{x} of the LP relaxation. Since there are only two constraints other than the bound constraints, at least

Algorithm $\mathcal{A}(\Pi)$ for $\Pi(S_1, S_2, O)$.

Let $\mu_{min} = \min_{j \in S_1} \mu_j$. Initialize $N_l = \lfloor \log_{1+\epsilon} \mu_{min} \rfloor$, $N_h = \lceil \log_{1+\epsilon} B \rceil$, $x_s = 0$, $P_s \leftarrow 0$.

1. For $t = N_l, \dots, N_h$,

(a) Let $\mu = (1 + \epsilon)^t$ and let $\tilde{x}(\mu)$ be a basic optimal solution for $\Pi(S_1, S_2, O, \mu)$.

(b) Using Lemma 3.1 find an integral solution $\hat{x}(\mu)$ such that

$$\sum_{j=1}^n p_j \cdot \hat{x}(\mu)_j \geq \sum_{j=1}^n p_j \cdot \tilde{x}(\mu)_j - 2\epsilon \cdot O$$

(c) If $P_s < \sum_{j \in S_1} p_j \hat{x}(\mu)_j + \sum_{j \in S_2} p_j$, then

$$\begin{aligned} x_s &\leftarrow \hat{x}(\mu) \\ P_s &\leftarrow \sum_{j \in S_1} p_j \hat{x}(\mu)_j + \sum_{j \in S_2} p_j \end{aligned}$$

2. Return the solution x_s .

Figure 2: Algorithm $\mathcal{A}(\Pi)$ for $\Pi(S_1, S_2, O)$

($|S_1| - 2$) bound constraints must be tight for \tilde{x} . Therefore, at least ($|S_1| - 2$) variables out of $|S_1|$ variables are integral in the basic optimal solution. Let $j_1, j_2 \in S$ such that $\tilde{x}_{j_1}, \tilde{x}_{j_2}$ are fractional. We know that $p_j \leq \epsilon \cdot O$ for all $j \in S_1$. Consider the following solution,

$$\hat{x}_j = \begin{cases} \tilde{x}_j, & j \in S_1, j \neq j_1, j_2, \\ 0, & \text{otherwise.} \end{cases}$$

Clearly, $\sum_{j \in S_1} p_j \hat{x}_j \geq \sum_{j \in S_1} p_j \tilde{x}_j - 2\epsilon \cdot O$. Therefore, we obtain an integral solution \hat{x} such that,

$$\sum_{j \in S_1} p_j \hat{x}_j + \sum_{j \in S_2} p_j \geq P^* - 2\epsilon \cdot O.$$

In the following lemma we show that for an appropriately chosen value of O and μ and subsets $S_1, S_2 \subset [n]$, the problem $\Pi(S_1, S_2, O, \mu)$ has optimal profit at least $\text{OPT}/(1 - \epsilon)$. ■

Lemma 3.2 *Let S^* be the set of items selected by an optimal solution and let $\text{OPT} = \sum_{i \in S^*} p_i$. Consider l , such that,*

$$(1 + \epsilon)^{l-1} \leq \text{OPT} < (1 + \epsilon)^l.$$

Let $O = (1 + \epsilon)^l$ and let,

$$S_\epsilon = \{i \in [n] \mid p_i \geq \epsilon \cdot O\}, S_1 = [n] \setminus S_\epsilon, S_2 = S_\epsilon \cap S^*.$$

Then the optimal profit for the problem $\Pi(S_1, S_2, O)$ is at least $\text{OPT}/(1 + \epsilon)$.

Proof: Let $\mu^* = \sum_{j \in S^*} \mu_j$, $\mu^1 = \sum_{j \in S_1 \cap S^*} \mu_j$, $\mu^2 = \sum_{j \in S_2} \mu_j$ and let k be such that $(1 + \epsilon)^{k-1} \leq \mu^1 < (1 + \epsilon)^k$. Let $\beta = (1 + \epsilon)^{k-1}$ and we consider the problem $\Pi(S_1, S_2, O, \beta)$. Consider the following fractional solution, \tilde{x} for $\Pi(S_1, S_2, O, \beta)$:

$$\tilde{x}_j = \begin{cases} \frac{1}{1+\epsilon}, & j \in S_1 \cap S^*, \\ 0, & \text{otherwise.} \end{cases}$$

We show that \tilde{x} is a feasible fractional solution for $\Pi(S_1, S_2, O, (1 + \epsilon)^{k-1})$ as:

$$\sum_{j \in S_1} \mu_j \tilde{x}_j = \sum_{j \in S_1 \cap S^*} \mu_j \cdot \frac{1}{\epsilon} = \frac{\mu^1}{1 + \epsilon} \leq \beta,$$

which implies that \tilde{x} satisfies (3.2). Also,

$$\begin{aligned} (\phi^{-1}(p))^2 \cdot \left(\sum_{j \in S_1} \sigma_j^2 \tilde{x}_j \right) &= (\phi^{-1}(p))^2 \cdot \left(\sum_{j \in S_1 \cap S^*} \sigma_j^2 \cdot \frac{1}{1 + \epsilon} \right) \\ &\leq \frac{(B - \mu^*)^2 - (\phi^{-1}(p))^2 \cdot \left(\sum_{j \in S_2} \sigma_j^2 \right)}{1 + \epsilon} \end{aligned} \quad (3.4)$$

$$\begin{aligned} &= \frac{(B - \mu^1 - \mu^2)^2 - (\phi^{-1}(p))^2 \cdot \left(\sum_{j \in S_2} \sigma_j^2 \right)}{1 + \epsilon} \\ &\leq \frac{(B - \beta - \mu^2)^2 - (\phi^{-1}(p))^2 \cdot \left(\sum_{j \in S_2} \sigma_j^2 \right)}{1 + \epsilon} \quad (3.5) \\ &< (B - \beta - \mu^2)^2 - (\phi^{-1}(p))^2 \cdot \left(\sum_{j \in S_2} \sigma_j^2 \right), \end{aligned}$$

where (3.4) follows as $S^* = (S_1 \cap S^*) \cup S_2$ is an optimal solution and thus, satisfies

$$(\phi^{-1}(p))^2 \cdot \left(\sum_{j \in S^*} \sigma_j^2 \right) \leq (B - \mu^*)^2,$$

and (3.5) follows as $\beta \leq \mu^1$. This implies that \tilde{x} satisfies (3.3) as well and thus, is a feasible solution for $\Pi(S_1, S_2, O, \beta)$. The profit achieved by the fractional solution \tilde{x} is:

$$\begin{aligned} \sum_{j \in S_1} p_j \tilde{x}_j + \sum_{j \in S_2} p_j &= \sum_{j \in S_1 \cap S^*} \frac{p_j}{1 + \epsilon} + \sum_{j \in S_2} p_j \\ &> \frac{\sum_{j \in S_1 \cap S^*} p_j + \sum_{j \in S_2} p_j}{1 + \epsilon} \\ &= \frac{\text{OPT}}{1 + \epsilon}, \end{aligned}$$

where the last equality follows because $S^* = (S_1 \cap S^*) \cup S_2$. Therefore, the optimal value for the problem $\Pi(S_1, S_2, O)$ is at least $\text{OPT}/(1 + \epsilon)$. \blacksquare

We now show that for any $\epsilon > 0$, the algorithm \mathcal{A} gives a $(1 - 3\epsilon)$ -approximation for the chance-constrained knapsack problem in running time $\tilde{O}(n^{1/\epsilon}/\epsilon^2)$.

Theorem 3.3 *Given $\epsilon > 0$, there is a polynomial time algorithm that gives a $(1 - 3\epsilon)$ -approximation for the chance constrained knapsack problem. Furthermore, the running time of \mathcal{A} is,*

$$O\left(\frac{\log(B/\mu_m) \cdot \log(P/p_m) \cdot n^{\frac{1}{\epsilon}}}{\epsilon^2}\right),$$

where $P = \sum_{j=1}^n p_j$, $p_m = \min\{p_j \mid j = 1, \dots, n\}$, $\mu_m = \min\{\mu_j \mid j = 1, \dots, n\}$.

Proof: Let OPT denote an optimal solution and let S^* be the set of items selected in OPT . Consider l , such that, $(1 + \epsilon)^{l-1} \leq \text{OPT} < (1 + \epsilon)^l$ and let $O = (1 + \epsilon)^l$. Let $S_\epsilon = \{i \in [n] \mid p_i \geq \epsilon \cdot O\}$, $S_1 = [n] \setminus S_\epsilon$ and $S_2 = S \cap S^*$. Note that the algorithm \mathcal{A} considers the guess O for the optimal value. Also, since $|S_2| < \frac{1}{\epsilon}$ the subproblem $\Pi(S_1, S_2, O)$ is considered as one of the subproblems in the algorithm \mathcal{A} . Let $\mu^1 = \sum_{j \in S_1 \cap S^*} \mu_j$. Consider k such that $(1 + \epsilon)^{k-1} \leq \mu^1 < (1 + \epsilon)^k$ and let $\beta = (1 + \epsilon)^{k-1}$. Clearly, the subproblem $\Pi(S_1, S_2, O, \beta)$ is considered in the algorithm $\mathcal{A}(\Pi)$ while solving $\Pi(S_1, S_2, O)$. From Lemma 3.2, we know that the optimal profit for the subproblem $\Pi(S_1, S_2, O, \beta)$ is at least $\frac{\text{OPT}}{1+\epsilon}$. Furthermore, using Lemma 3.1 we can find a set of items \hat{S} for the problem $\Pi(S_1, S_2, O, \beta)$ such that,

$$\begin{aligned} \sum_{j \in \hat{S}} p_j &\geq \frac{\text{OPT}}{1 + \epsilon} - 2\epsilon \cdot O \\ &\geq \left(\frac{1}{1 + \epsilon} - 2\epsilon \right) \cdot \text{OPT} \\ &\geq (1 - 3\epsilon) \cdot \text{OPT} \end{aligned}$$

Therefore, the algorithm \mathcal{A} finds an integral solution that has profit at least $(1 - 3\epsilon) \cdot \text{OPT}$. To bound the running time of \mathcal{A} , note that we consider $O(\log(P/p_m)/\epsilon)$ different choices of the optimal profit value O . Also, we consider $O(n^{1/\epsilon})$ choices of the set of items S for the subproblem Π for each choice of O . Furthermore, in the subroutine $\mathcal{A}(\Pi)$, we solve $O(\log(B/\mu_m)/\epsilon)$ different sub-problems for solving $\Pi(S_1, S_2, O)$ for given subsets $S_1, S_2 \subset [n]$ and a choice for optimal profit O . Therefore, the total running time of \mathcal{A} is,

$$O\left(\frac{\log(B/\mu_m) \cdot \log(P/p_m) \cdot n^{\frac{1}{\epsilon}}}{\epsilon^2}\right).$$

■

References

- [1] S.P. Boyd and L. Vandenberghe. *Convex Optimization*. Cambridge University Press, 2004.
- [2] R.L. Carraway, R.L. Schmidt, and L.R. Weatherford. An Algorithm for Maximizing Target Achievement in the Stochastic Knapsack Problem with Normal Returns. *Naval Research Logistics*, 40:161–161, 1993.
- [3] B. Dean, M. Goemans, and J. Vondrak. The benefit of adaptivity: approximating the stochastic knapsack problem. In *Proceedings of the 45th IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 208–217, 2004.
- [4] A. Goel and P. Indyk. Stochastic Load Balancing and Related Problems. In *Annual Symposium On Foundations Of Computer Science*, volume 40, pages 579–586, 1999.
- [5] MI HENIG. Risk criteria in a stochastic knapsack problem. *Operations research*, 38(5):820–825, 1990.
- [6] J. Kleinberg, Y. Rabani, and É. Tardos. Allocating bandwidth for bursty connections. In *Proceedings of the twenty-ninth annual ACM symposium on Theory of computing*, pages 664–673. ACM New York, NY, USA, 1997.
- [7] A.J. Kleywegt and J.D. Papastavrou. The Dynamic and Stochastic Knapsack Problem with Random Sized Items. *Operations Research*, 49(1):26, 2001.
- [8] J.D. Papastavrou, S. Rajagopalan, and A.J. Kleywegt. The Dynamic and Stochastic Knapsack Problem with Deadlines. *Management Science*, 42:1706–1718, 1996.