Learning Algorithms for Separable Approximations of Discrete Stochastic Optimization Problems

Warren B. Powell
Department of Operations Research and Financial Engineering,
Princeton University, Princeton, NJ 08544

Andrzej Ruszczyński
Department of Management Science and Information Systems
Rutgers University
94 Rockefeller Rd, Piscataway, NJ 08854

Huseyin Topaloglu
Department of Operations Research and Industrial Engineering
Cornell University
Ithaca, NY

March 10, 2004

Abstract

We propose the use of sequences of separable, piecewise linear approximations for solving nondifferentiable stochastic optimization problems. The approximations are constructed adaptively using a combination of stochastic subgradient information and possibly sample information on the objective function itself. We prove the convergence of several versions of such methods when the objective function is separable and has integer break points, and we illustrate their behavior on numerical examples. We then demonstrate the performance on nonseparable problems that arise in the context of two-stage stochastic programming problems, and demonstrate that these techniques provide near optimal solutions with a very fast rate of convergence compared to other solution techniques.

Keywords: Stochastic programming
MSC 2000 classification codes: 90C15 (Stochastic programming)

1 Introduction

We consider the following stochastic programming problem

$$\max_{x \in X} \mathbb{E} f(x, \omega), \tag{1}$$

where $f : \mathbb{R}^n \times \Omega \to \mathbb{R}$, $(\Omega, \mathcal{F}, \mathbb{P})$ is a probability space, and $\mathbb{E}$ denotes the expected value. We assume that the function f is concave with respect to the first argument, and such that

$$F(x) = \mathbb{E} f(x, \omega)$$

is finite for every $x \in X$. We assume that for almost all ω, all $i = 1, \ldots, m$, and for all feasible integer values of x_j for $j \neq i$, the function $f(x_1, \ldots, x_{i-1}, \cdot, x_{i+1}, \ldots, x_n, \omega)$ is piecewise linear with integer break points. We also assume that the set X is closed, and

$$X \subseteq \{x \in \mathbb{R}^n : 0 \leq x_i \leq M_i,\ i = 1, \ldots, n\}.$$

Problems of this type arise in a variety of resource allocation problems. In the car distribution problem of railroads, for example, planners often face the problem of having to reposition empty cars before a customer order has been realized. Truckload carriers have to assign drivers to loads before knowing the opportunities available to the driver at the destination. The air mobility command needs logic for their simulators that will reposition aircraft back to locations where they *might* be needed, before actual demands are known.

All of these problems can be modeled as multistage stochastic linear programming problems. Experimental research (see, for example, Godfrey & Powell (2001)) has demonstrated that the recourse function for these problems can be well approximated by sequences of piecewise linear, separable functions. These scalar functions can be approximated by techniques that use sample gradients, where the challenge is to maintain the concavity of the approximations. As a result of the need to use Monte Carlo samples, it is necessary to introduce steps which maintain the concavity after *every* update. This paper explores a class of such techniques and establishes conditions under which these approximations converge to the true function.

An important case is what is referred to as the two-stage stochastic program:

$$\begin{aligned} \max \quad & \langle c, x \rangle + \mathbb{E}\ Q(x, \omega) \\ \text{subject to:} \quad & Ax = b, \\ & x \geq 0, \end{aligned} \tag{2}$$

where

$$\begin{aligned} Q(x, \omega) = \ & \max \langle q, y \rangle \\ \text{subject to:} \quad & Wy = h(\omega) - Tx, \\ & y \geq 0. \end{aligned}$$

A special case of (2) is the two-stage stochastic program with network recourse, where the second stage problem is the min-cost network flow problem. If h takes integer values a.s.,

then this problem is a special case of (1). Another special case occurs when f in (1) is a separable function of the form

$$f(x,\omega) = \sum_{i=1}^{n} f_i(x_i,\omega). \tag{3}$$

This is the form most often taken in the context of classical resource allocation problems (Righter (1989)) which involve the allocation of resources to independent activities subject to a common budget constraint.

In order to obtain the optimal solution to (1), one can consider building sequential approximations of F, say F^k. If the sequence $\{F^k\}$ converges to F in an appropriate sense then we can claim to have a procedure to solve (1). Alternatively we may solve optimization problems of the form $x^k \in \arg\max_{x\in X} F^k(x)$, constructed in such a way that the sequence $\{x^k\}$ converges to $x^* \in \arg\max_{x\in X} F(x)$. In this case the sequence of functions $\{F^k\}$ does not necessarily converge to F, but $\{x^k\}$ may converge to an point. Of critical importance in practical applications is also the speed of convergence, a question that we treat on an experimental basis.

For our class of applications, it is relatively easy, for a given x^k, to sample an elementary event ω^k and to calculate $f_i(x_i^k,\omega^k)$. Moreover, it is also easy to obtain information about the slope of $f_i(\cdot,\omega^k)$ at x_i^k:

$$v_i^k = f_i(x_i^k,\omega^k) - f_i(x_i^k - 1,\omega^k). \tag{4}$$

On the other hand, it is difficult to obtain the exact values of $\bar{f}_i(x) = \mathbb{E}\{f_i(x,\omega)\}$, since it involves the calculation of the expected value.

Example 1 Let x_i denote the amount of resource allocated to activity i, where $i = 1,\ldots,n$. These amounts have to be chosen from the set

$$X = \Big\{x \in \mathbb{R}^n : x_i \in \{0,1,\ldots,M_i\},\ i = 1,\ldots,n,\ \sum_{i=1}^{n} x_i \le b\Big\}. \tag{5}$$

For each activity i there is a nonnegative integer random variable D_i representing the demand. The reward associated with activity i is defined as in the newsvendor problem:

$$f_i(x_i,D_i) = q_i \min(x_i,D_i) - c_i x_i,$$

where $q_i > c_i > 0$. Our objective is to allocate the resources in such a way that the expected reward, $F(x) = \mathbb{E}\sum_{i=1}^{n} f_i(x_i,D_i)$, is maximized, subject to the constraint $x \in X$. Porteus (1990) provides a thorough review of the newsvendor problem in the context of stochastic inventory models. The optimal solution of a single newsvendor problem can be expressed analytically, but this requires knowing the distribution of demand. An extensive literature has evolved to solve what is known as the censored newsvendor problem (where you only see the amount sold, not the actual demand). This literature (see, for example, Ding et al. (2002)) requires assuming a parametric form for the demand distribution. The algorithm

provided in this paper does not require any information about the demand distribution or a set of demand observations from which the demand distribution can empirically be constructed.

In such a problem we may sample a demand realization $D^k = (D_1^k, \ldots, D_n^k)$ and calculate

$$v_i^k = \begin{cases} q_i - c_i & \text{if } x_i^k \leq D_i^k, \\ -c_i & \text{if } x_i^k > D_i^k. \end{cases} \quad i = 1, \ldots, n.$$

We can also calculate the right slope estimate

$$v_{i+}^k = \begin{cases} q_i - c_i & \text{if } x_i^k < D_i^k, \\ -c_i & \text{if } x_i^k \geq D_i^k, \end{cases} \quad i = 1, \ldots, n.$$

The information our algorithm uses is $\{x^k\}$, $\{v^k\}$ and $\{v_+^k\}$, where x^k, v^k and v_+^k respectively denote the vectors $(x_1^k, \ldots, x_n^k)$, $(v_1^k, \ldots, v_n^k)$ and $(v_{1+}^k, \ldots, v_{n+}^k)$. In order to compute any v_i^k or v_{i+}^k, we only need to know if $x_i^k \leq D_i^k$. For the newsvendor problem, this translates into knowing whether the newsvendor has sold all the newspapers or not, rather than observing the exact value of the demand random variable. This is exactly the same situation addressed in the censored newsvendor problem. Therefore, the algorithm presented in this paper provides an asymptotically optimal solution to the censored newsvendor problem without requiring any particular form for the demand distribution.

Similar estimates can be generated in a slightly more complicated case, with the reward associated with activity i defined as

$$f_i(x_i, D_i) = q_i(\min(x_i, D_i)) - c_i(x_i),$$

where $q_i(\cdot)$ is a concave piecewise linear function, and $c_i(\cdot)$ is a convex piecewise linear function, both with break points at $0, 1, \ldots, M_i$.

There exists a wealth of numerical methods for stochastic programming problems. The first group are scenario methods, in which a sufficiently rich sample $\omega^1, \ldots, \omega^N$ is drawn from the space Ω, and the expectation is approximated by the sample average:

$$F^N(x) = \frac{1}{N} \sum_{\nu=1}^{N} f(x, \omega^\nu).$$

A discussion of these approaches can be found in Shapiro & Homem-De-Mello (2000), Korf & Wets (2001) and Kleywegt et al. (2002).

The second group of methods are *stochastic subgradient methods*, which use the fact that the random vector v^k in (4) satisfies the relation

$$\mathbb{E}\{v^k | x^k\} \in \partial F(x^k),$$

where ∂F is the subdifferential of F (understood as the negative of the subdifferential of the convex function $-F$). Stochastic subgradient algorithms depend on updates of the form:

$$x^{k+1} = x^k + \alpha_k v^k.$$

These methods make very many small updates to the current approximation by using the stochastic subgradients as directions, achieving convergence almost surely to the optimal point. A treatment of these methods can be found in Ermoliev (1988). Constraints can be treated by projection, feasible direction techniques Ruszczyński (1980) or by recursive linearization Ruszczyński (1987).

Finally, Benders decomposition can be used to solve two-stage stochastic linear programs by approximating the recourse function $\mathbb{E}\,Q(x,\omega)$ with a series of cutting planes (Van Slyke & Wets (1969), Ruszczyński (1986), Higle & Sen (1991), Ruszczyński (2003), Chen & Powell (1999)).

Our problem class is motivated by problems that require integer solutions. The introduction of uncertainty often has the effect of destroying the natural integer structure of many problems. All of the three classes of techniques mentioned above destroy the natural integrality either by how the recourse function is approximated (scenario methods and Benders decomposition) or the nature of the algorithm itself (stochastic subgradient methods). The method proposed in this paper handles integrality requirements very easily (Laporte & Louveaux (1993) shows how integrality can be incorporated into Benders decomposition).

We propose to solve these problems by adaptively estimating, using sample subgradient information, sequences of *separable* approximations which are piecewise linear, concave, and have integer break points. We use the information gathered in iterations $1,\ldots,k$ to construct models $f_i^k(\cdot)$ of the expected value functions $\bar{f}_i(\cdot) = \mathbb{E}\{f_i(\cdot,\omega)\}$, $i = 1,\ldots,n$. The next approximation to the solution is given by:

$$x^k \in \arg\max_{x\in X} \sum_{i=1}^{n} f_i^k(x_i). \tag{6}$$

An associated learning step provides information employed to update the models f_i^k. Such an approach is already known to be optimal if the objective function is continuously differentiable (Culioli & Cohen (1990), Cheung & Powell (2000)), but there is no comparable result for nondifferentiable problems. While the relation of stochastic approximation type methods and learning is well known (see, e.g. Kushner & Yin (1997)), the use of the structure (separability and concavity) allows here for the construction of particularly efficient methods.

Our solution strategy extends a line of research in stochastic resource allocation using separable approximations. This problem class has been most widely studied using the framework of two-stage stochastic programs with network recourse (Wallace (1986), Wallace (1987) and Birge & Wallace (1988)). Independently, separable, piecewise linear approximations have been proposed for discrete resource allocation problems that arise in the context of fleet management (Powell (1986), Powell (1987) and Powell (1988)). Frantzeskakis & Powell (1990) suggests a static, piecewise linear separable approximation for specially structured tree problems, a result that is generalized in Powell & Cheung (1994) and applied to multistage resource allocation problems in Cheung & Powell (1996). These methods, however, were not adaptive which limited the quality of the solution. Powell & Carvalho (1998) provided an adaptive learning algorithm based on linear approximations, which was extended

in Godfrey & Powell (2001) to an adaptive, piecewise linear approximation based on the "CAVE" algorithm. The CAVE algorithm provided exceptionally good experimental performance, but offered no provable results. Wallace (1987) introduces a piecewise linear upper bound for networks, a result that is generalized in Birge & Wallace (1988) for stochastic programs.

In this paper, we introduce and formally study the use of sequences of piecewise linear, separable approximations as a strategy for solving nondifferentiable stochastic optimization problems. As a byproduct, we produce a fast algorithm for problems such as two stage stochastic programs with network recourse, a topic that was first studied in depth by Wallace (1986). We establish several important convergence results for the special case of separable objective functions, and show experimentally that the algorithm provides near-optimal, and often optimal, solutions for problems when the objective function is not separable, as would be the case with two-stage stochastic programs. Furthermore, the observed speed of convergence is much faster than techniques such as Benders decomposition, especially for higher dimensional problems.

The paper is divided into two parts. Sections 2-6 deal exclusively with problems where the original objective function $F(x, \omega)$ is separable. While this problem class enjoys its own sets of applications (for example, in a variety of budget allocation problems), our interest in this special problem class arises primarily because we are able to prove some important convergence results. Section 2 presents the basic algorithm for learning piecewise linear, concave approximations (while maintaining concavity after every update), and proves convergence to the real function assuming that all points are sampled infinitely often. Section 3 provides a variation of the algorithm that combines gradient information with sample information on the function itself. In practical applications, we cannot generally guarantee that we will sample all points infinitely often, and this is not necessary to find the optimal solution. Section 4 proves convergence when we only sample the points x^k generated by equation (6). Section 5 shows how a certain projection required to maintain concavity can be implemented, Section 6 provides the results of a series of experiments that investigate the rate of convergence of variations of the algorithm.

The second part of the paper, given in section 7, focuses on nonseparable problems that arise in the context of two-stage stochastic programs. We cannot guarantee that our algorithm will produce the optimal solution for two-stage problems, but we show that the right separable approximation can produce the optimal solution, and use this to develop a bound on our result. Numerical comparisons with Benders decomposition, which is optimal for non-integer versions of these problems, indicate that our approach may provide much faster convergence and optimal or very near-optimal results.

2 Learning Concave Functions of One Variable

We start from the description and analysis of the basic learning algorithm for a concave piecewise linear function of one variable $\bar{f} : [0, M] \rightarrow \mathbb{R}$. We assume that $\bar{f}$ is linear on the

intervals $[s-1, s]$, $s = 1, \ldots, M$. Let

$$\bar{v}_s = \bar{f}(s) - \bar{f}(s-1), \quad s = 1, \ldots, M.$$

Let us note that the knowledge of the vector $\bar{v} = (\bar{v}_1, \ldots, \bar{v}_M)$ allows us to reconstruct $\bar{f}(x)$, $x \in [0, M]$, except for the constant term $\bar{f}(0)$:

$$\bar{f}(x) = \bar{f}(0) + \sum_{s=1}^{l} \bar{v}_s + \bar{v}_{l+1}(x - l),$$

where l is such that $l \leq x < l + 1$. In the context of our problem (1) with the objective function (3), $\bar{f}(\cdot)$ represents the expected value of a coordinate function, $\mathbb{E}\{f_i(\cdot)\}$.

The main idea of the algorithm is to recursively update a random vector v^k taking values in $\mathbb{R}^M$, $k = 1, 2, \ldots$, in order to achieve convergence of v^k to $\bar{v}$ (in some stochastic sense). We still denote by $(\Omega, \mathcal{F}, \mathbb{P})$ the probability space on which this sequence is defined.

Let us note that by the concavity of $\bar{f}$ the vector $\bar{v}$ has nonincreasing components:

$$\bar{v}_{s+1} \leq \bar{v}_s, \quad s = 1, \ldots, M-1. \tag{7}$$

We shall at first assume that there exists a constant B such that

$$\bar{v}_1 \leq B, \quad \bar{v}_M \geq -B. \tag{8}$$

Clearly, the set V of vectors satisfying (7)–(8) is convex and closed. We shall ensure that all our approximate slopes v^k are elements of V as well. To this end we shall employ the operation of orthogonal projection on V

$$\Pi_V(z) = \arg\min\{\|v - z\|^2 : v \in V\}. \tag{9}$$

We show in section 5 that for the set V defined by (7) such a projection can be calculated in an easy way.

Our learning algorithm, which is called the separable, projective approximation routine (SPAR), is given in figure 1.

At this moment we shall not specify the way in which s^k is defined, except that s^k is a random variable. Specific conditions on it will be formulated later. We use $\mathcal{F}_k$ to denote the σ-subalgebra generated by $v^1, \ldots, v^k, s^1, \ldots, s^{k-1}$. We denote

$$p_s^k = \mathbb{P}\{s^k = s \,|\, \mathcal{F}_k\}, \quad s = 1, \ldots, M.$$

The stepsizes α_k employed at step 3 may also be random, but must be $\mathcal{F}_k$-measurable.

Let us denote by ξ^k the random vector with the components

$$\xi_s^k = \begin{cases} -\eta^k + v_s^k & \text{if } s = s^k, \\ 0 & \text{otherwise.} \end{cases}$$

Step 0 Set $v^1 \in V$, $k = 1$.

Step 1 Sample $s^k \in \{1, \ldots, M\}$.

Step 2 Observe a random variable η^k such that

$$\mathbb{E}\{\eta^k \,|\, v^1, \ldots, v^k; s^1, \ldots, s^k\} = \bar{v}_{s^k}, \text{ a.s..} \tag{10}$$

Step 3 Calculate the vector $z^k \in \mathbb{R}^M$ as follows

$$z_s^k = \begin{cases} (1-\alpha_k)v_s^k + \alpha_k \eta^k & \text{if } s = s^k, \\ v_s^k & \text{otherwise,} \end{cases} \tag{11}$$

where $\alpha_k \in (0, 1]$.

Step 4 Calculate $v^{k+1} = \Pi_V(z^k)$, increase k by one and go to step 1.

Figure 1: Separable, Projective Approximation Routine (SPAR)

We can now rewrite the method compactly as

$$v^{k+1} = \Pi_V(v^k - \alpha_k \xi^k), \quad k = 1, 2, \ldots.$$

It follows from (10) that

$$\mathbb{E}\{\xi_s^k \,|\, \mathcal{F}_k\} = p_s^k(v_s^k - \bar{v}_s), \quad s = 1, \ldots, M.$$

Thus

$$\mathbb{E}\{\xi^k \,|\, \mathcal{F}_k\} = P^k(v^k - \bar{v}), \quad P^k = \operatorname{diag}(p_s^k)_{s=1}^M. \tag{12}$$

We assume that there exists a constant C such that for all k

$$\mathbb{E}\{\|\eta^k\|^2 \,|\, v^1, \ldots, v^k; s^1, \ldots, s^k\} \leq C, \text{ a.s.}, \quad k = 1, 2, \ldots. \tag{13}$$

We also assume that

$$\sum_{k=1}^{\infty} \alpha_k = \infty, \text{ a.s.}, \tag{14}$$

$$\sum_{k=1}^{\infty} \mathbb{E}\alpha_k^2 < \infty, \tag{15}$$

$$\liminf_{k\to\infty} p_s^k > 0, \text{ a.s.}, \quad s = 1, \ldots, M. \tag{16}$$

.

Theorem 1 *Assume (10) and (13)–(16). Then SPAR generates a sequence* $\{v^k\}$ *such that* $v^k \to \bar{v}$ *a.s..*

Proof. Our proof is standard, but we present it here in order to derive some useful inequalities that will be applied later. By the nonexpansiveness of the projection $\Pi_V(\cdot)$,

$$\|v^{k+1} - \bar{v}\|^2 \leq \|z^k - \bar{v}\|^2 = \|v^k - \bar{v}\|^2 - 2\alpha_k \langle v^k - \bar{v}, \xi^k \rangle + \alpha_k^2 \|\xi^k\|^2.$$

We add and subtract the term $2\alpha_k \langle v^k - \bar{v}, P^k(v^k - \bar{v}) \rangle$ to obtain:

$$\begin{aligned} \|v^{k+1} - \bar{v}\|^2 \leq & \|v^k - \bar{v}\|^2 - 2\alpha_k \langle v^k - \bar{v}, P^k(v^k - \bar{v}) \rangle \\ & - 2\alpha_k \langle v^k - \bar{v}, \xi^k - P^k(v^k - \bar{v}) \rangle + \alpha_k^2 \|\xi^k\|^2. \end{aligned} \tag{17}$$

Let us consider the sequence

$$S_m = \sum_{k=1}^{m} \alpha_k^2 \|\xi^k\|^2, \quad m = 1, 2 \ldots,$$

and let $S_0 = 0$. By the boundedness of V, and by (13) there exists a constant C_1 such that $\mathbb{E}\{\alpha_m^2 \|\xi^m\|^2 \,|\, \mathcal{F}_m\} \leq C_1 \alpha_m^2$ a.s., for all m. Therefore, in view of the $\mathcal{F}_m$-measurability of α_m,

$$S_{m-1} \leq \mathbb{E}\{S_m \,|\, \mathcal{F}_m\} \leq S_{m-1} + C_1 \alpha_m^2, \quad m = 1, 2, \ldots.$$

Taking the expected value we obtain that $\mathbb{E}\{S_m\} \leq \mathbb{E}\{S_{m-1}\} + C_1 \mathbb{E}\{\alpha_m^2\}$ for all m, and thus

$$\mathbb{E}\{S_m\} \leq C_1 \mathbb{E}\Big\{ \sum_{k=1}^{m} \alpha_k^2 \Big\}.$$

The last two displayed relations and assumption (15) imply that the sequence $\{S_m\}$ is a submartingale, which is convergent a.s., by virtue of (Doob 1953, Thm. 4.1).

Consider now the series

$$U_m = \sum_{k=1}^{m} \alpha_k \langle v^k - \bar{v}, \xi^k - P^k(v^k - \bar{v}) \rangle, \quad m = 1, 2, \ldots, \tag{18}$$

and let $U_0 = 0$. By (12), $\mathbb{E}\{U_m \,|\, \mathcal{F}_m\} = U_{m-1}$, $m = 1, 2, \ldots$, and thus the sequence $\{U_m\}$ is a martingale. We can write equation (18) as:

$$U_m = U_{m-1} + \alpha_m \langle v^m - \bar{v}, \xi^m - P^m(v^m - \bar{v}) \rangle.$$

Squaring both sides and taking the expectation yields:

$$\begin{aligned} \mathbb{E}\{U_m^2 \,|\, \mathcal{F}_m\} &= U_{m-1}^2 + \mathbb{E}\Big\{ \big[\alpha_m \langle v^m - \bar{v}, \xi^m - P^m(v^m - \bar{v}) \rangle\big]^2 \,\Big|\, \mathcal{F}_m \Big\} \\ &\quad + \mathbb{E}\Big\{ U_{m-1} \Big(\alpha_m \langle v^m - \bar{v}, \xi^m - P^m(v^m - \bar{v}) \rangle \Big) \,\Big|\, \mathcal{F}_m \Big\} \qquad (19) \\ &= U_{m-1}^2 + \mathbb{E}\Big\{ \big[\alpha_m \langle v^m - \bar{v}, \xi^m - P^m(v^m - \bar{v}) \rangle\big]^2 \,\Big|\, \mathcal{F}_m \Big\}. \qquad (20) \end{aligned}$$

where the last term in equation (19) is zero due to (12). By the boundedness of V, and by (13), there exists a constant C_2 such that

$$\mathbb{E}\Big\{ \big[\alpha_m \langle v^m - \bar{v}, \xi^m - P^m(v^m - \bar{v})\rangle\big]^2 \,\Big|\, \mathcal{F}_m \Big\} \le C_2 \alpha_m^2, \quad m = 1, 2, \ldots. \tag{21}$$

Equations (20) and (21) yield:

$$\mathbb{E}\{U_m^2 \,|\, \mathcal{F}_m\} \le U_{m-1}^2 + C_2 \alpha_m^2, \quad m = 1, 2, \ldots.$$

Taking the expected value we conclude that

$$\mathbb{E}\{U_m^2\} \le C_2 \mathbb{E}\Big\{ \sum_{k=1}^{m} \alpha_k^2 \Big\}, \quad m = 1, 2, \ldots.$$

Assumption (15) implies that the martingale $\{U_m\}$ is convergent a.s., by virtue of (Doob 1953, Thm. 4.1). Therefore (17) may be rewritten as

$$\|v^{k+1} - \bar{v}\|^2 \le \|v^k - \bar{v}\|^2 - 2\alpha_k \langle v^k - \bar{v}, P^k(v^k - \bar{v})\rangle + A_k. \tag{22}$$

where $A_k = (S_k - S_{k-1}) - (U_k - U_{k-1})$ and $\sum_{k=1}^{\infty} A_k$ is finite a.s. This implies that

$$\|v^{k+1} - \bar{v}\|^2 + \sum_{j=k+1}^{\infty} A_j \le \|v^k - \bar{v}\|^2 + \sum_{j=k}^{\infty} A_j, \quad k = 1, 2, \ldots.$$

The sequence $\|v^k - \bar{v}\|^2 + \sum_{j=k}^{\infty} A_j$, $k = 1, 2, \ldots$ is nonincreasing and bounded from below a.s., hence convergent. Thus the sequence $\{\|v^k - \bar{v}\|^2\}$ is convergent a.s. From (22) we get

$$\sum_{k=1}^{\infty} \alpha_k \langle v^k - \bar{v}, P^k(v^k - \bar{v})\rangle < \infty, \text{ a.s..}$$

Using (14) and (16) we deduce that a.s. there must exist an infinite subset of indices $\mathcal{K} \subseteq \mathbb{N}$ and a subsequence $\{v^k\}$, $k \in \mathcal{K}$, such that $v^k \to \bar{v}$ for $k \in \mathcal{K}$. Since the sequence of distances $\|v^k - \bar{v}\|^2$ is convergent, the entire sequence $\{v^k\}$ converges to $\bar{v}$. □

If we remove inequalities (8) from the definition of V, only small technical changes are needed to ensure convergence a.s. Instead of the steps $\alpha_k \xi^k$ we need to use normalized steps $\alpha_k \gamma_k \xi^k$, where the normalizing coefficients have the form:

$$\gamma_k = (\max(\|v^k\|, B))^{-1},$$

for some large constant B. We first prove that both martingales $\{S_m\}$ and $\{U_m\}$ converge, due to the damping by the γ_k's. Then the corresponding version of (22) yields the boundedness of $\{v^k\}$ a.s. Consequently, the normalizing coefficients are bounded away from 0, a.s., and the remaining part of the analysis goes through, as well. In our further considerations we shall still assume, though, that inequalities (8) are present in the definition of V, in order to avoid unnecessary notational complications associated with the normalizing coefficients γ_k.

In many applications, at a given point $s^k \in \{1, \ldots, M-1\}$ we can observe *two* random variables: η^k satisfying (10), and η^k_+ such that

$$\mathbb{E}\{\eta^k_+ \,|\, v^1, \ldots, v^k; s^1, \ldots, s^k\} = \bar{v}_{s^k+1} \tag{23}$$

and

$$\mathbb{E}\{\|\eta^k_+\|^2 \,|\, v^1, \ldots, v^k; s^1, \ldots, s^k\} \leq C, \quad k = 1, 2, \ldots. \tag{24}$$

This was illustrated in example 1.

Our algorithm can be easily adapted to this case, too. The only difference is step 3, where we use both random observations, whenever they are available:

$$z^k_s = \begin{cases} (1-\alpha_k)v^k_s + \alpha_k \eta^k & \text{if } s = s^k, \\ (1-\alpha_k)v^k_s + \alpha_k \eta^k_+ & \text{if } s^k < M \text{ and } s = s^k + 1, \\ v^k_s & \text{otherwise,} \end{cases} \tag{25}$$

The analysis of this version of the method is similar to the basic case. We define

$$\xi^k_s = \begin{cases} -\eta^k + v^k_s & \text{if } s = s^k, \\ -\eta^k_+ + v^k_s & \text{if } s^k < M \text{ and } s = s^k + 1, \\ 0 & \text{otherwise.} \end{cases}$$

It follows from (10) and (23) that

$$\mathbb{E}\{\xi^k_s | \mathcal{F}_k\} = \begin{cases} p^k_s (v^k_s - \bar{v}_s) & \text{if } s = 1, \\ (p^k_s + p^k_{s-1})(v^k_s - \bar{v}_s) & \text{if } 1 < s \leq M. \end{cases} \tag{26}$$

Therefore, after replacing the coefficients p^k_s by

$$\bar{p}^k_s = \begin{cases} p^k_s & \text{if } s = 1 \\ p^k_s + p^k_{s-1} & \text{if } 1 < s \leq M, \end{cases}$$

we can reduce this version of the method to the basic case analyzed earlier.

3 Using Objective Value Observations

In the applications that we have in mind, our observations provide us with more information than just the estimate of the slope of the objective function at s^k. We also observe the value of the objective function at s^k corresponding to some outcome ω^k. We denote this by

$$\theta^k = f(s^k, \omega^k).$$

Step 0 Set $v^1 \in V$, $k = 1$.

Step 1 Sample $s^k \in \{1, \ldots, M\}$.

Step 2 Observe random variables η^k and θ^k satisfying (10) and (27).

Step 3 Calculate the vector $z^k \in \mathbb{R}^M$ as follows

$$z_s^k = \begin{cases} v_s^k + \alpha_k \Big(\theta^k - \sum\limits_{i=1}^{s^k} v_i^k\Big) & \text{for } s = 1, \ldots, s^k - 1, \\ (1-\alpha_k) v_s^k + \alpha_k \eta^k + \alpha_k \Big(\theta^k - \sum\limits_{i=1}^{s^k} v_i^k\Big) & \text{for } s = s^k, \\ v_s^k & \text{otherwise,} \end{cases} \tag{28}$$

where $\alpha_k \in (0, 1]$.

Step 4 Calculate $v^{k+1} = \Pi_V(z^k)$, increase k by one and go to step 1.

Figure 2: SPAR with objective function updates (SPAR-Obj)

Usually, we know the value $\bar{f}(0)$ and with no loss of generality we assume $\bar{f}(0) = 0$. We also assume that θ^k satisfies

$$\mathbb{E}\{\theta^k | v^1, \ldots, v^k; s^1, \ldots, s^k\} = \bar{f}(s^k) = \sum_{i=1}^{s^k} \bar{v}_i, \text{ a.s.} \tag{27}$$

(27) is trivially satisfied when ω^k is independent of $v^1, \ldots, v^k; s^1, \ldots, s^k$ but this does not necessarily have to be the case. We can now use $\{\theta^k\}$ to facilitate the convergence to $\bar{v}$. The algorithm is described in figure 2.

We additionally assume that there exists a constant C such that

$$\mathbb{E}\{(\theta^k)^2 \,|\, v^1, \ldots, v^k; s^1, \ldots, s^k\} \le C, \quad k = 1, 2, \ldots. \tag{29}$$

We have a result similar to theorem 1.

Theorem 2 *Assume (10), (13)–(16) and (27), (29). Then SPAR-Obj generates a sequence* $\{v^k\}$ *such that* $v^k \to \bar{v}$ *a.s..*

Proof. Let us calculate the conditional expectation of the vector $\xi^k = (v^k - z^k)/\alpha_k$. Directly from (10) and (27) we have

$$\Big(\mathbb{E}\{\xi^k \,|\, \mathcal{F}_k\}\Big)_s = \begin{cases} \sum\limits_{i=1}^{s^k} (v_i^k - \bar{v}_i^k) & \text{for } s = 1, \ldots, s^k - 1, \\ v_{s^k}^k - \bar{v}_{s^k}^k + \sum\limits_{i=1}^{s^k} (v_i^k - \bar{v}_i^k) & \text{for } s = s^k, \\ 0 & \text{for } s = s^k + 1, \ldots, M. \end{cases}$$

Therefore

$$
\begin{aligned}
(\mathbb{E}\{\xi^k|\mathcal{F}_k\})_j &= p_j^k(v_j^k - \bar{v}_j) + \sum_{s=j}^{M} p_s^k \sum_{i=1}^{s}(v_i^k - \bar{v}_i^k) \\
&= \sum_{i=1}^{j-1}\Big(\sum_{s=j}^{M} p_s^k\Big)(v_i^k - \bar{v}_i^k) + \Big(p_j^k + \sum_{s=j}^{M} p_s^k\Big)(v_j^k - \bar{v}_j^k) \\
&\quad + \sum_{i=j+1}^{M}(\sum_{s=i}^{M} p_s^k)(v_i^k - \bar{v}_i^k).
\end{aligned}
$$

Consider the matrix W^k of dimension $M \times M$ with the entries

$$
w_{ij}^k = \begin{cases} \sum_{s=j}^{M} p_s^k & \text{if } i \le j, \\ \sum_{s=i}^{M} p_s^k & \text{if } i > j. \end{cases}
$$

The last two relations yield:

$$
\mathbb{E}\{\xi^k|\mathcal{F}_k\} = (P^k + W^k)(v^k - \bar{v}^k), \tag{30}
$$

where $P^k = \operatorname{diag}(p_j^k)_{j=1}^M$. We have

$$
W^k = \sum_{l=1}^{M} W_l^k,
$$

where each of the matrices W_l^k has entries

$$
(W_l^k)_{ij} = \begin{cases} p_l^k & \text{if } i, j \le l, \\ 0 & \text{otherwise.} \end{cases}
$$

Each W_l^k is positive semidefinite, and thus W^k is positive semidefinite, too. Now we can proceed as in the proof of theorem 1. We have an inequality similar to (17):

$$
\begin{aligned}
\|v^{k+1} - \bar{v}\|^2 \le \|v^k - \bar{v}\|^2 &- 2\alpha_k\langle v^k - \bar{v}, (P^k + W^k)(v^k - \bar{v})\rangle \\
&- 2\alpha_k\langle v^k - \bar{v}, \xi^k - (P^k + W^k)(v^k - \bar{v})\rangle + \alpha_k^2\|\xi^k\|^2.
\end{aligned}
$$

Using (30) and (15) we can rewrite it as

$$
\|v^{k+1} - \bar{v}\|^2 \le \|v^k - \bar{v}\|^2 - 2\alpha_k\langle v^k - \bar{v}, (P^k + W^k)(v^k - \bar{v})\rangle + A_k.
$$

where $\sum_{k=1}^{\infty} A_k$ is finite a.s. Since W^k is positive semidefinite, we can just omit it in the above inequality and obtain (22) again. The rest of the proof is the same as before. □

SPAR-Obj can be modified to assign different weights to the components of the direction associated with the observations η^k and θ^k. In particular, we may set at step 3

$$
z_s^k = \begin{cases} v_s^k + \frac{\alpha_k}{\rho_k}\Big(\theta^k - \sum_{i=1}^{s^k} v_i^k\Big) & \text{for } s = 1, \ldots, s^k - 1, \\ (1-\alpha_k)v_s^k + \alpha_k\eta^k + \frac{\alpha_k}{\rho_k}\Big(\theta^k - \sum_{i=1}^{s^k} v_i^k\Big) & \text{for } s = s^k, \\ v_s^k & \text{otherwise.} \end{cases} \tag{31}
$$

where $0 < \rho^{\min} \leq \rho_k \leq \rho^{\max}$, and ρ_k is $\mathcal{F}_k$-measurable. The analysis of this version is identical to the proof of theorem 2. Our numerical experiments reported in section 6 indicate that the additional scaling in (31) is useful.

4 Multidimensional Problems and Learning While Optimizing

Let us now return to problem (1). If the next observation point $s^k = (s_1^k, \ldots, s_n^k)$ is sampled at random, and if for each coordinate i the probabilities

$$p_{is}^k = \mathbb{P}\{s_i^k = s \,|\, \mathcal{F}_k\}, \quad s = 1, \ldots, M,$$

satisfy assumption (16), then all results of the preceding two sections apply component-wise.

The situation is different if we generate the next observation point s^k by solving the approximate problem (6), that is

$$s_i^k = x_i^k, \quad i = 1, \ldots, n, \quad k = 1, 2, \ldots.$$

If the solution x^k is not unique, we choose it at random from the set of optimal solutions of (6). For each coordinate function f_i we observe two random variables: η_i^k satisfying (10), and η_{i+}^k satisfying (23). Then we update the left and right slopes for each function f_i according to (25), and the iteration continues. In this way we define the sequences $\{v_i^k\}$, $k = 1, 2 \ldots$ of estimates of the slopes of f_i, $i = 1, \ldots, n$, and a sequence x^k of the solutions of approximate models (6).

This algorithm is well-defined if the approximate problems (6) have integer solutions for all concave piecewise linear functions f_i having integer break points. This is true, for example, for models having alternative network representations, as those discussed in Powell & Topaloglu (2003).

Note that in the previous two sections our convergence proofs depend on the assumption that $\liminf_{k\to\infty} p_{is}^k > 0$ a.s. for $s = 1, \ldots, M$, $i = 1, \ldots, n$. However, when s^k is selected as $s^k = \arg\max_{x\in X} \sum_{i=1}^n \bar{f}_i^k(x_i)$, this assumption may not satisfied. In this section, we show that even with this new choice of s^k, the sequence $\{s^k\}$ converges to an optimal solution of $\max_{x\in X} \sum_{i=1}^n \bar{f}_i(x_i)$, provided a certain stability condition is satisfied.

Let us note that for a concave, piecewise linear and separable function

$$F(x) = \sum_{i=1}^n f_i(x_i), \tag{32}$$

where each f_i is defined as

$$f_i(x_i) = \sum_{s=1}^l v_s + v_{l+1}(x - l) \tag{33}$$

with an integer l such that $l \le x < l+1$, the subdifferential of F at an *integer* point x is given by

$$\partial F(x) = [v_{1,x_1+1}, v_{1,x_1}] \times [v_{2,x_2+1}, v_{2,x_2}] \times \cdots \times [v_{n,x_n+1}, v_{n,x_n}].$$

Under the Slater constraint qualification, the necessary and sufficient condition of optimality for problem

$$\max_{x \in X} F(x),$$

where X is a convex closed set, has the form

$$0 \in \partial F(x) - N(x), \tag{34}$$

with $N(x)$ being the normal cone to X at x. An optimal point $\hat{x}$ is called *stable* if it satisfies

$$0 \in \operatorname{int} [\partial F(\hat{x}) - N(\hat{x})]. \tag{35}$$

It can be seen directly from conditions (34) and (35) that a stable point $\hat{x}$ is also a solution to a perturbed problem

$$\max_{x \in X} \tilde{F}(x),$$

provided that $\operatorname{dist}(\partial F(\hat{x}), \partial \tilde{F}(\hat{x})) < \epsilon$ and ϵ is a sufficiently small positive number.

Clearly, the solutions x^k of our approximate problems (6) satisfy condition (34) for the approximate functions f^k constructed by the method. Then, by passing to the limit, we can conclude that each accumulation point (x^*, v^*) of the sequence $\{(x^k, v^k)\}$ satisfies the condition

$$0 \in \partial F^*(x^*) - N(x^*),$$

with F^* constructed from v^* as in (32)–(33). We shall show that if such an accumulation point satisfies the condition of stability, it is optimal for the original problem.

Theorem 3 *Assume that for each $i = 1, \ldots, n$ the conditions (10), (13)–(15) and (23)–(24) are satisfied. If an accumulation point (x^*, v^*) of the sequence $\{(x^k, v^k)\}$ generated by the algorithm, satisfies the stability condition:*

$$0 \in \operatorname{int} [\partial F^*(x^*) - N(x^*)], \tag{36}$$

then with probability one x^* is an optimal solution of (1).

Proof. Let us observe that relation (26) holds for each coordinate i. Therefore inequality (22) is true for each coordinate i:

$$\|v_i^{k+1} - \bar{v}_i\|^2 \le \|v_i^k - \bar{v}_i\|^2 - 2\alpha_k \langle v_i^k - \bar{v}_i, P_i^k (v_i^k - \bar{v}_i)\rangle + A_{ik}. \tag{37}$$

The matrix P_i^k, which is $\mathcal{F}_k$-measurable, is a nonnegative diagonal matrix with positive entries corresponding to the i-th coordinates of possible solutions to (6). Proceeding exactly as in the proof of theorem 1 we conclude that the series $\sum_{k=1}^{\infty} A_{ik}$ is convergent a.s. Furthermore, the sequence $\{\|v_i^k - \bar{v}_i\|\}$ is convergent a.s., for every $i = 1, \ldots, n$.

Our proof will analyze properties of sample paths of the random sequence $\{(v^k, x^k)\}$ for all elementary events $\omega \in \Omega \setminus \Omega_0$, where Ω_0 is a null set. It will become clear in the course of the proof what this null set is.

Let us fix $\omega \in \Omega$ and consider a convergent subsequence $\{(v^k(\omega), x^k(\omega))\}$, $k \in \mathcal{K}(\omega)$, where $\mathcal{K}(\omega) \subseteq \mathbb{N}$ is some infinite set of indices. Let us denote by (v^*, x^*) the limit of this subsequence. This limit depends on ω too, but we shall omit the argument ω to simplify notation.

If the stability condition holds, then there exists $\epsilon > 0$ such that for all iterations k for which $|v_{i,x_i^*}^k(\omega) - v_{i,x_i^*}^*| \leq \epsilon$, $i = 1, \ldots, n$, the solution x^k of the approximate problem (6) is equal to x^*. Then the coefficients $p_{i,s}^k$ are equal to 1 for $s = x_i^*$ and $s = x_i^* + 1$, and are zero otherwise, for each function i. Let us fix an arbitrary i and focus our attention on the points $s = x_i^*$. Inequality (37) implies:

$$\|v_i^{k+1}(\omega) - \bar{v}_i\|^2 \leq \|v_i^k(\omega) - \bar{v}_i\|^2 - 2\alpha_k(\omega)(v_{i,x_i^*}^k(\omega) - \bar{v}_{i,x_i^*})^2 + A_{ik}(\omega). \tag{38}$$

The series $\sum_k A_{ik}$ is convergent a.s. Let $k \in \mathcal{K}(\omega)$ be large enough so that $|v_{i,x_i^*}^k(\omega) - v_{i,x_i^*}^*| < \epsilon/2$. Consider $j \geq k$ such that

$$|v_{i,x_i^*}^j(\omega) - v_{i,x_i^*}^*| \leq \epsilon \quad \text{for all} \quad i = 1, \ldots, n. \tag{39}$$

Let us suppose that the ith coordinate of the limit point is not optimal, i.e.,

$$v_{i,x_i^*}^* \neq \bar{v}_{i,x_i^*}. \tag{40}$$

We shall prove that it leads to a contradiction. The remaining part of our proof has three stages.

Stage 1

We can always choose a sufficiently small $\epsilon > 0$ such that $|v_{i,x_i^*}^* - \bar{v}_{i,x_i^*}| > 2\epsilon$. Then for the iterations j satisfying (39) we have $|v_{i,x_i^*}^j(\omega) - \bar{v}_{i,x_i^*}| > \epsilon$, and inequality (38) implies:

$$\|v_i^{j+1}(\omega) - \bar{v}_i\|^2 \leq \|v_i^j(\omega) - \bar{v}_i\|^2 - 2\alpha_j(\omega)\epsilon^2 + A_{ij}(\omega). \tag{41}$$

The series $\sum_j \alpha_j(\omega)$ is divergent and the series $\sum_j A_{ij}(\omega)$ is convergent, unless ω is in a certain null set. If the set of consecutive $j \geq k$ for which condition (39) holds was infinite, inequality (41) would lead to a contradiction. Therefore, for all $k \in \mathcal{K}(\omega)$ and all sufficiently small $\epsilon > 0$, the random index

$$l(k, \epsilon, \omega) = \inf\{j \geq k : \max_{1 \leq i \leq n} |v_{i,x_i^*}^j(\omega) - v_{i,x_i^*}^*| > \epsilon\}$$

is finite.

Stage 2

We shall prove that the sum of stepsizes between $k \in \mathcal{K}(\omega)$ and $l(k,\epsilon,\omega)-1$ is at least of order ϵ, if k is large enough. By the definition of $l(k,\epsilon,\omega)$ we have, for some i, $|v_{i,x_i^*}^{l(k,\epsilon,\omega)}(\omega) - v_{i,x_i^*}^*| > \epsilon$. Since $v^k(\omega) \to v^*$, $k \in \mathcal{K}(\omega)$, we also have $\|v_i^{l(k,\epsilon,\omega)} - v_i^k(\omega)\| > \epsilon/2$ for all sufficiently large $k \in \mathcal{K}(\omega)$. Thus

$$\sum_{j=k}^{l(k,\epsilon,\omega)-1} \alpha_j(\omega)\|\xi_i^j(\omega)\| > \epsilon/2. \tag{42}$$

Let us observe that conditions (13) and (15) imply that for each i the random series

$$\sum_{k=1}^{\infty} \alpha_k\big(\|\xi_i^k\| - \mathbb{E}\{\|\xi_i^k\| \mid \mathcal{F}_k\}\big)$$

is a convergent martingale. Therefore, unless ω is in a certain null set,

$$\sum_{j=k}^{l(k,\epsilon,\omega)-1} \alpha_j(\omega)\|\xi_i^j(\omega)\| = \sum_{j=k}^{l(k,\epsilon,\omega)-1} \alpha_j(\omega)\kappa_{ij}(\omega) + \sigma_{ik}(\omega),$$

where $\kappa_{ij} = \mathbb{E}\{\|\xi_i^j\| \mid \mathcal{F}_j\}$ and $\sigma_{ik}(\omega) = \sum_{j=k}^{l(k,\epsilon,\omega)-1} \alpha_j(\omega)(\|\xi_i^j(\omega)\| - \kappa_{ij}(\omega)) \to 0$, as $k \to \infty$, $k \in \mathcal{K}(\omega)$. This combined with (42) implies that for all sufficiently large $k \in \mathcal{K}(\omega)$

$$\sum_{j=k}^{l(k,\epsilon,\omega)-1} \alpha_j(\omega)\kappa_{ij}(\omega) \geq \epsilon/3.$$

From assumption (13) it follows that there exists a constant C such that $\kappa_{ij}(\omega) \leq C$ for all i and j. Using this in the last displayed inequality we obtain

$$\sum_{j=k}^{l(k,\epsilon,\omega)-1} \alpha_j(\omega) \geq \frac{\epsilon}{3C}, \tag{43}$$

for all sufficiently small $\epsilon > 0$ and all sufficiently large $k \in \mathcal{K}(\omega)$.

Stage 3

Summing (41) from k to $l(k,\epsilon,\omega)-1$, letting $\delta = 1/3C$ and combining with (43) gives, for some coordinate i:

$$\|v_i^{l(k,\epsilon,\omega)}(\omega) - \bar{v}_i\|^2 \leq \|v_i^k(\omega) - \bar{v}_i\|^2 - 2\delta\epsilon^3 + \sum_{j=k}^{l(k,\epsilon,\omega)-1} A_{ij}(\omega).$$

Let $\Delta(\omega)$ be the limit of the entire sequence $\{\|v_i^j(\omega) - \bar{v}_i\|^2\}$, as $j \to \infty$, whose existence has been established at the beginning of the proof. Passing to the limit with $k \to \infty$, $k \in \mathcal{K}(\omega)$, and using the fact that $\sum_{j=k}^{\infty} A_{ij}(\omega) \to 0$, as $k \to \infty$, we obtain

$$\Delta(\omega) \leq \Delta(\omega) - 2\delta\epsilon^3,$$

a contradiction. Therefore our assumption (40) must be false, and we have

$$v^*_{i,x_i^*} = \bar{v}_{i,x_i^*} \quad \text{for all} \quad i = 1, \ldots, n. \tag{44}$$

Inequality (38) is also true with x_i^k replaced by $x_i^k + 1$ (if $x_i^k < M_i$). We can thus apply the same argument to prove

$$v^*_{i,x_i^*+1} = \bar{v}_{i,x_i^*+1} \quad \text{for all} \quad i = 1, \ldots, n. \tag{45}$$

For $x_i^* = M_i$ we take the convention that $v^*_{i,x_i^*+1} = \bar{v}_{i,x_i^*+1} = -\infty$. Consequently,

$$\partial F(x^*) = \partial F^*(x^*)$$

and the point x^* is optimal for (1). □

The assumptions of theorem 3 are stronger than those of theorems 1 and 2. However, its result is much stronger too. For a general closed convex set X, it may be very costly to devise a sampling scheme for $\{s^k\}$ satisfying $\liminf_{k\to\infty} p_{is}^k > 0$ a.s. for $s = 1, \ldots, M$, $i = 1, \ldots, n$. Theorem 3 saves us from devising such a sampling scheme and lets us pick s^k by simply solving an optimization problem. The stability assumption (36) is difficult to verify *a priori*, but it is very easy to check *a posteriori*, when the accumulation point x^* and the approximate function F^* have been identified.

In a similar way (and under identical assumptions) we can prove the convergence of the version that uses function value estimates.

Theorem 4 *Assume (10), (13)–(15), (23)–(24), and (27), (29). If an accumulation point (x^*, v^*) of the sequence $\{(x^k, v^k)\}$ generated by SPAR-Obj, satisfies the stability condition (36) then with probability one x^* is an optimal solution of (1).*

The proof is almost a *verbatim* copy of the proof of theorem 3, with the modifications as in theorem 2.

5 Projection

Let us now describe the way the projection $v = \Pi_V(z)$ can be calculated. Clearly, v is the solution to the quadratic programming problem

$$\min \frac{1}{2}\|v - z\|^2 \tag{46}$$

$$\text{subject to: } v_{s+1} - v_s \leq 0, \quad s = 0, \ldots, M, \tag{47}$$

where for uniformity we denote $v_0 = B$, $v_{M+1} = -B$. Associating with (47) Lagrange multipliers $\lambda_s \geq 0$, $s = 0, \ldots, M$, we obtain the necessary and sufficient optimality conditions:

$$v_s = z_s + \lambda_s - \lambda_{s-1}, \quad s = 1, \ldots, M, \tag{48}$$

$$\lambda_s(v_{s+1} - v_s) = 0, \quad s = 0, \ldots, M. \tag{49}$$

If $i_1, \ldots, i_2$ is a sequence of coordinates such that

$$v_{i_1-1} > v_{i_1} = v_{i_1+1} = \cdots = c = \cdots = v_{i_2-1} = v_{i_2} > v_{i_2+1},$$

then adding the equations (48) from i_1 to i_2 yields

$$c = \frac{1}{i_2 - i_1 + 1} \sum_{s=i_1}^{i_2} z_s.$$

If $i_1 = 1$, then c is the minimum of the above average and B, and for $i_2 = M$ the maximum of $-B$ and this average has to be taken.

The second useful observation is that $v^k \in V$ and z^k computed by (11) differs from v^k in just one coordinate. If $z^k \notin V$, one of two cases must occur: either $z^k_{s^k-1} < z^k_{s^k}$, or $z^k_{s^k+1} > z^k_{s^k}$.

If $z^k_{s^k-1} < z^k_{s^k}$, we search for the largest $1 < i \leq s^k$ for which

$$z^k_{i-1} \geq \frac{1}{s^k - i + 1} \sum_{s=i}^{s^k} z^k_s. \tag{50}$$

If such i cannot be found we set $i = 1$. Then we calculate

$$c = \frac{1}{s^k - i + 1} \sum_{s=i}^{s^k} z^k_s \tag{51}$$

and set

$$v^{k+1}_j = \min(B, c), \quad j = i, \ldots, s^k. \tag{52}$$

We have $\lambda_0 = \max(0, c - B)$, and

$$\lambda_s = \begin{cases} 0 & s = 1, \ldots, i-1, \\ \lambda_{s-1} + z_s - v_s & s = i, \ldots, s^k - 1, \\ 0 & s = s^k, \ldots, M. \end{cases}$$

It is straightforward to verify that the solution found and the above Lagrange multipliers satisfy conditions (48)–(49).

The procedure in the case when $z^k_{s^k} < z^k_{s^k+1}$ is symmetrical: it is the same procedure applied to the graph of z rotated by π.

Let us now consider the method which employs *two* random variables at each iteration, with step 3 as in (25). Then both z_{s^k} and z_{s^k+1} may differ from v_{s^k} and v_{s^k+1} (although we still have $z_{s^k} > z_{s^k+1}$). We shall show that algorithm (50)–(52) can easily be adapted to this case.

Suppose that $z_{s^k} > z_{s^k-1}$. We apply (50)–(51) to compute the candidate value for c. Now two cases may occur.

Case 1. If $c \geq z_{s^k+1}$, we may apply (52). We can now focus on the points to the right: $z_{s^k+1}, z_{s^k+2}, \ldots$. We apply the symmetric analogue of (50)–(52) to these values, and the projection is accomplished.

Case 2. If $c < z_{s^k+1}$, the value of c is not correct. We need to include z_{s^k+1} into the averaging procedure. Thus we repeat (50)–(52) but with s^k replaced by $s^k + 1$ in (50), although we still search for $1 < i \leq s^k$. After this we apply (52) for $j = i, \ldots, s^k + 1$.

If $z_{s^k+1} < z_{s^k+2}$ the situation is symmetrical to the one discussed above (after rotating the graph of z by π), and an analogous procedure can be applied.

6 Experiments for Separable Problems

To illustrate the behavior of the methods discussed, we consider the problem in example 1:

$$\max_{x \in X} \mathbb{E} \sum_{i=1}^{n} f_i(x_i, D_i), \text{ where } X \text{ is given in (5) and } f_i(x_i, D_i) = q_i \min(x_i, D_i) - c_i x_i.$$

Clearly, both SPAR and SPAR-Obj can be applied to this problem componentwise: We approximate $\mathbb{E} \sum_{i=1}^{n} f_i(x_i, D_i)$ by $\sum_{i=1}^{n} f_i^k(x_i)$ at iteration k. For a given resource allocation $s^k = (s_1^k, \ldots, s_n^k)$ among n activities and a sampled demand realization $D^k = (D_1^k, \ldots, D_n^k)$, we can separately apply the updates of SPAR and SPAR-Obj for each component $i = 1, \ldots, n$. In the description below s_i^k plays the role of s^k for the i-th component, ρ_{ki} the role of ρ_k, etc. We compare the following methods:

- SPAR - This is the basic learning algorithm with projection.
- SPAR-Obj(a) with $\rho_{ki} = s_i^k$ - This uses objective function estimates to help with the learning, using weights of $\frac{1}{s^k}$ for the objective function estimates.
- SPAR-Obj(b) with $\rho_k = M_i s_i^k$ - Same as above, but with a much smaller weight on the objective function.
- The Leveling Method of Topaloglu & Powell (2003) - This algorithm maintains concavity by forcing slopes that violate an updated estimate to be no larger (to the left) or no smaller (to the right) than the most recently updated cell. This algorithm has been shown to be convergent.

- The CAVE Algorithm of Godfrey & Powell (2001) - This was the first algorithm suggested for adaptively estimating piecewise linear functions while maintaining concavity.

In the first series of experiments, the random variable s^k is generated from the uniform distribution over the rectangle $[1, M_1] \times \cdots \times [1, M_n]$. We assume that each component of the demand variable D^k is independent and Poisson distributed. Having sampled s_i^k and D_i^k, we compute the left-hand slope of $f_i(\cdot, D_i^k)$ at s_i^k as

$$\eta_i^k = \begin{cases} q_i - c_i & \text{if } s_i^k \leq D_i^k \\ -c_i & \text{if } s_i^k > D_i^k. \end{cases}$$

Having obtained this slope information and the value of s_i^k for iteration k, we can obtain the approximation f_i^{k+1} using any of the methods mentioned above. We call this a *learning step*, since in this iteration we are exploring the slope of the function at randomly sampled points.

In order to estimate the quality of an approximation we find $\bar{x}^k \in \arg\max_{x \in X} \sum_{i=1}^n f_i^k(x_i)$ and compute $\mathbb{E} \sum_{i=1}^n f_i(\bar{x_i}^k, D_i)$. This gives us an idea about the average actual performance of the solution given by the approximation $\sum_{i=1}^n f_i^k$.

For each method and coordinate (activity) the sequence of step sizes is $\alpha_k = 20/(40+k)$. We take the number of coordinates to be $n = 90$. M_i ranges between 20 and 40, c_i ranges between 0.6 and 1.4 for different 90 activities. D_i^k is truncated-Poisson distributed with mean ranging between 9 and 21 for $i = 1, \ldots, 90$. Finally $q_i = 2$ for all $i = 1, \ldots, n$ and $b = 950$. We run each method 50 times using 100 demand realizations at each run and figure 3(a) presents the averages over these 50 runs.

We see that our basic learning method, SPAR, performs very well. Its quality can be slightly improved by using objective function estimates as in SPAR-Obj, but the weight associated to them must be significantly smaller than the weight associated with the slope observations, as the comparison of versions SPAR-Obj(a) and SPAR-Obj(b) shows.

The second series of experiments differ from the first by the method we utilized to generate the random variables s^k. Now, s^k is chosen to be the maximizer of $\sum_{i=1}^n f_i^k$, as discussed in section 4. We call this an *optimizing step* since this step involves maximizing the current approximation as opposed to selecting s^k randomly over the interval $[1, M_1] \times \cdots \times [1, M_n]$. This version concentrates its efforts around the maximizers of the approximations, and one might expect that it has a potential of being more efficient. The results are collected in figure 3(b).

Comparing figures 3(a) and 3(b), all optimizing methods perform worse than the corresponding learning method at the early iterations. Only after about 30-40 iterations did the versions with optimizing steps take the lead over their learning counterparts. However, the tail performance of the optimizing methods is much better.

Several conclusions can be drawn from our experiments. First, the application of the projection operator facilitates the convergence. It provides an update to a range of values on the basis of the observation obtained at one value. Second, learning is useful especially

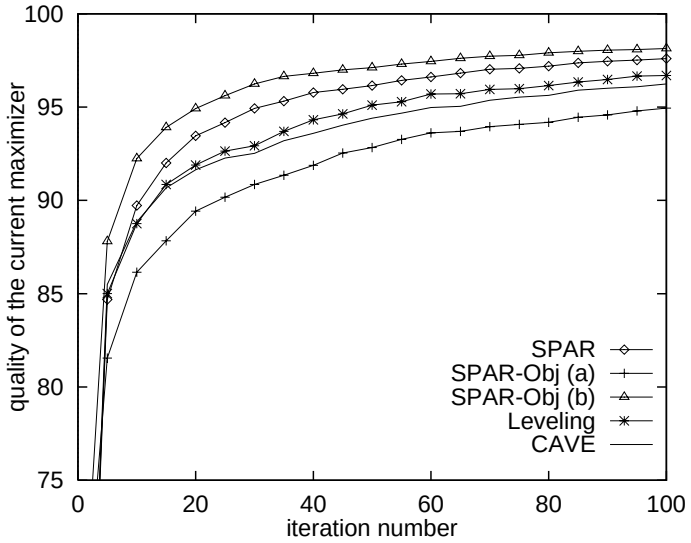

3a: Comparison of methods with learning steps.

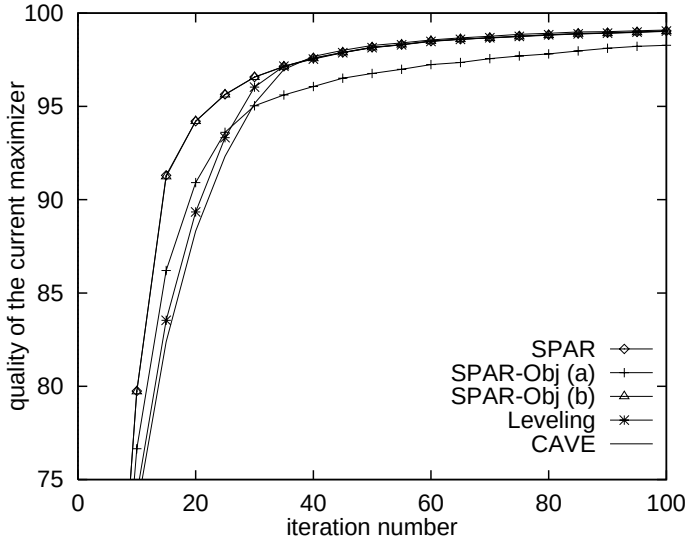

3b: Comparison of methods with optimizing steps.

Figure 3: Comparison of methods with learning (a) and optimizing (b)

at earlier iterations. Instead of trying to shoot at the optimal point for the current model, it is better to collect information at randomly selected points from time to time. Third, the indirect use of noisy objective values to correct the subgradients has a positive effect on the convergence, provided that the weight of the additional modification is small. Finally, the learning-based methods provide good approximations to the solutions at early iterations, which makes them attractive candidates for problems where the cost of one experiment is high.

7 Application to Nonseparable Resource Allocation Problems

We now turn our attention to nonseparable problems that arise in the context of two-stage stochastic resource allocation problems. Section 7.1 shows that there exists a separable approximation that will produce an optimal first-stage solution, although there is no guarantee that our algorithm will find this approximation. We also provide a bound on the solution provided by our algorithm. Then, section 7.2 compares our algorithm, which produces integer solutions, to the solutions produced by several variations of Benders decomposition.

7.1 Outline of the method

Let us start from the following observation. Consider the problem

$$\max_{x \in X} \; F(x), \tag{53}$$

where $X \subset \mathbb{R}^n$ is a closed convex set, and $F : \mathbb{R}^n \to \mathbb{R}$ is a concave function. Suppose that $\hat{x}$ is the optimal solution of the problem, and that F is subdifferentiable at $\hat{x}$. Let us construct a concave separable approximation of F at $\hat{x}$ in the form

$$\bar{F}^{\hat{x}}(x) = F(\hat{x}) + \sum_{i=1}^{n} \bar{f}_i^{\hat{x}}(x_i),$$

where

$$\bar{f}_i^{\hat{x}}(x_i) = \begin{cases} F'(\hat{x}, e_i)(x_i - \hat{x}_i) & \text{if } x_i \geq \hat{x}_i, \\ F'(\hat{x}, -e_i)(\hat{x}_i - x_i) & \text{if } x_i \leq \hat{x}_i. \end{cases}$$

In the formula above $F'(\hat{x_i}, d)$ denotes the directional derivative of F at $\hat{x}$ in direction d, and e_i is the i-th unit vector in $\mathbb{R}^n$. We use $\hat{x}$ as the superscript of $\bar{F}$ to stress that the approximation is constructed at $\hat{x}$.

The point $\hat{x}$ is also the solution of the deterministic approximate problem

$$\max_{x \in X} \; \bar{F}^{\hat{x}}(x). \tag{54}$$

Indeed, each direction d can be represented as

$$d = \sum_{i=1}^{n} d_i e_i = \sum_{i=1}^{n} (d_i)_+ e_i + \sum_{i=1}^{n} (d_i)_- (-e_i).$$

Since the directional derivative is a concave positively homogeneous function, we have

$$F'(\hat{x}, d) \geq \sum_{i=1}^{n} (d_i)_+ F'(\hat{x}, e_i) + \sum_{i=1}^{n} (d_i)_- F'(\hat{x}, -e_i) = \left[\bar{F}^{\hat{x}}\right]'(\hat{x}, d). \tag{55}$$

By the optimality of $\hat{x}$, the directional derivative $F'(\hat{x}, d)$ of F at $\hat{x}$ in any direction $d \in \text{cone}\,(X - \hat{x})$ is non-positive. Therefore $\left[\bar{F}^{\hat{x}}\right]'(\hat{x}, d) \leq 0$ for every feasible direction d, as required.

Consider an arbitrary point y, at which such a separable approximation $\bar{F}^y(x)$ has been constructed. We have just proved that if $y = \hat{x}$ then y is the solution of (54). The converse statement is not true, in general, and one can easily construct counter-examples in which the separable approximation $\bar{F}^y(\cdot)$ constructed at some point y achieves its maximum at y, but $F(\cdot)$ does not.

Then, however, we can derive an upper bound on the optimal value of $F(\cdot)$ in X as follows. If $\hat{x}$ is an optimal solution of problem (53), then

$$F(\hat{x}) - F(y) \leq F'(y, \hat{x} - y) \leq -F'(y, y - \hat{x}).$$

In the second inequality above we have used the concavity of $F'(y, \cdot)$. Inequality (55) can be developed at the point y instead of $\hat{x}$ and reads

$$F'(y, d) \geq \left[\bar{F}^y\right]'(y, d).$$

Setting $d = y - \hat{x}$ and combining the last two inequalities we obtain

$$F(\hat{x}) - F(y) \leq -F'(y, y - \hat{x}) \leq -\left[\bar{F}^y\right]'(y, y - \hat{x}).$$

Thus the following bound on the difference between the optimal value of problem (53) and the value of F at y holds true:

$$F(\hat{x}) - F(y) \leq -\min_{x \in X} \left[\bar{F}^y\right]'(y, y - x). \tag{56}$$

The quantity at the right hand side can be easily calculated or estimated, given the current approximate solution y and the piecewise linear separable approximation $\bar{F}^y$.

For a general stochastic programming problem with a nonseparable recourse function, our methods do not necessarily converge to the optimal solution. Furthermore, our methods use *s*amples of the directional derivatives in the directions e_i, rather than exact values, so the error bound will be an estimate, as well.

For solving $\max_{x \in X} F(x)$, when F is a nonseparable function, our method proceeds as follows: At iteration k, our approximation is $\sum_{i=1}^{n} f_i^k(x_i)$, where each f_i^k is a one-dimensional, piecewise linear, concave function characterized by the slope vector $v_i^k = (v_{i1}^k, \ldots, v_{iM_i}^k)$ and M_i is an upper bound on the variable x_i. The point x^k is the maximizer of $\sum_{i=1}^{n} f_i^k(x_i)$. We update our approximation using the slopes gathered at x^k. The algorithm is described in figure 4.

Several remarks are in order. First, as a stopping criterion, one may choose to continue for a specified number of iterations or until $F(x^k)$ does not improve for a certain number of iterations.

Second, we note that there are two sources of error in this approximate procedure. The main one is the use of separable approximations, as discussed above. The second one is

Step 0 Set $v_i^1 \in V$ for all $i = 1, \ldots, n$, $k = 1$.

Step 1 Find $x^k \in \arg\max_{x \in X} \sum_{i=1}^n f_i^k(x_i)$, where each f_i^k is defined by the slope vector v_i^k.

Step 2 Observe a random variable η^k such that

$$\mathbb{E}\{\eta^k \,|\, v^1, \ldots, v^k; x^1, \ldots, x^k\} \in \partial F(x^k), \text{ a.s..}$$

Step 3 For each $i = 1, \ldots, n$ do the following:

Step 3.1 Calculate the vector $z_i^k \in \mathbb{R}^{M_i}$ as follows

$$z_{is}^k = \begin{cases} (1-\alpha_k)v_{is}^k + \alpha_k \eta_i^k & \text{if } s = x_i^k, \\ v_{is}^k & \text{otherwise.} \end{cases}$$

Step 3.2 Calculate $v_i^{k+1} = \Pi_V(z_i^k)$.

Step 4 Increase k by one and go to step 1.

Figure 4: The optimizing version of SPAR

the use of an arbitrary stochastic subgradient η^k rather than estimates of the forward and backward directional derivatives, as required in section 4. Nevertheless, the method performs remarkably well on a class of stochastic optimization problems that we discuss below.

Third, many stochastic programming problems lend themselves to compact state variables and the recourse functions in these problems have considerably fewer dimensions than the number of decision variables. For example, in (2), the recourse cost $\mathbb{E}Q(x, \omega)$ depends on Tx. If the dimension of Tx is less than x, by writing (2) as

$$\begin{aligned} \max \ & \langle c, x \rangle + \mathbb{E}\ Q(s, \omega) \\ \text{subject to:} \ & Ax = b, \\ & Tx - s = 0, \\ & x \geq 0, \end{aligned}$$

where

$$\begin{aligned} Q(s, \omega) = \ & \max \langle q, y \rangle \\ \text{subject to:} \ & Wy = h(\omega) - s, \\ & y \geq 0, \end{aligned}$$

and building separable approximations of $\mathbb{E}\, Q(s, \omega)$, we can decrease the number of required approximations.

In this case, the extension of the algorithm we presented above is straightforward. The dimension of s is denoted by n. In step 1, we set $(x^k, s^k) \in \arg\max_{(x,s) \in X} \langle c, x \rangle + \sum_{i=1}^n f_i^k(s_i)$, where X is the feasible set of the first stage problem. In step 2, η^k has to satisfy $\mathbb{E}\{\eta^k \,|\, v^1, \ldots, v^k;\ s^1, \ldots, s^k\} \in \partial \mathbb{E}\,\{Q(s^k, \omega)\}$ One can choose η^k to be the Lagrange

multipliers associated with the constraints of the second stage problem for a certain realization of ω, say ω^k and for a certain value of s^k. If ω^k is independent of all previously used realizations of ω, then $\mathbb{E}\{\eta^k \mid v^1, \ldots, v^k; s^1, \ldots, s^k\} \in \partial\mathbb{E}\{Q(s^k, \omega)\}$ is easily satisfied.

7.2 Numerical Illustration

We illustrate our method using the following two-stage stochastic programming problem:

$$
\begin{aligned}
\max \quad & \sum_{i\in\mathcal{I}} \sum_{j\in\mathcal{I}\cup\mathcal{C}} c_{ij}x_{ij} + \mathbb{E}Q(s, D) \\
\text{subject to:} \quad & \sum_{j\in\mathcal{I}\cup\mathcal{C}} x_{ij} \le p_i, \quad i \in \mathcal{I}, \\
& \sum_{i\in\mathcal{I}} x_{ij} - s_j = 0, \quad j \in \mathcal{I}\cup\mathcal{C}, \\
& x_{ij}, s_j \ge 0,
\end{aligned}
$$

where $Q(s, D)$ is the optimal value of the second stage problem:

$$
\begin{aligned}
\max \quad & \sum_{i\in\mathcal{I}\cup\mathcal{C}} \sum_{j\in\mathcal{C}} d_{ij}y_{ij} + \sum_{i\in\mathcal{C}} \sum_{l\in\mathcal{L}} r_i^l z_i^l \\
\text{subject to:} \quad & \sum_{j\in\mathcal{C}} y_{ij} \le s_i, \quad i \in \mathcal{I}\cup\mathcal{C}, \\
& \sum_{i\in\mathcal{I}\cup\mathcal{C}} y_{ij} - \sum_{l\in\mathcal{L}} z_j^l \ge 0, \quad j \in \mathcal{C}, \\
& z_j^l \le D_j^l, \quad l \in \mathcal{L}, \quad j \in \mathcal{C}, \\
& y_{ij}, z_j^l \ge 0.
\end{aligned}
$$

The problem above can be interpreted as follows: There is a set of production facilities (with warehouses) $\mathcal{I}$ and a set of customers $\mathcal{C}$. At the first stage, an amount x_{ij} is transported from production facility i to a warehouse or customer location j, before the demand realizations at the customer locations become known. After the realizations of the demand at the customer locations are observed, we move an amount y_{ij} from location i to customer location j. At each customer location we face different types of demands, indexed by $l \in \mathcal{L}$: D_i^l is the demand of type l at location j. We serve z_j^l units of demand of type l at location j; the excess demand, if any, is lost. The production capacity of facility i is denoted by p_i.

For the first stage costs, we set $c_{ij} = c_0 + c_1\delta_{ij}$, where δ_{ij} is the Euclidean distance between locations i and j, and c_0 can be interpreted as the unit production cost and c_1 is the transportation cost applied on a per mile basis. For the second stage costs, we have

$$
d_{ij} = \begin{cases} d_1\delta_{ij} & \text{if } i \in \mathcal{I} \text{ or } i = j \\ d_0 + d_1\delta_{ij} & \text{if } i \in \mathcal{C} \text{ and } i \ne j. \end{cases}
$$

d_0 represents the fixed charge for shipping a unit of the product from one customer location to another customer location, and d_1 is the per mile cost of transportation in the second stage. For every demand type l occurring in location i, we associate a revenue r_i^l. Our test problems differ in cost parameters and $|\mathcal{I} \bigcup \mathcal{C}|$, which determines the dimensionality of the recourse function. They all have 100 possible demand scenarios.

As a benchmark, we use three well-known Benders decomposition-based stochastic programming algorithms: L-shaped decomposition (LSD) (Van Slyke & Wets (1969)), stochastic decomposition (SD) (Higle & Sen (1991)), and cutting plane partial sampling (CUPPS) (Chen & Powell (1999)).

Our focus is on the rate of convergence, measured by the improvement in the objective function as the number of iterations or the CPU time increases. In order to measure the rate of convergence of different methods, we ran each algorithm for 25, 100, 500, 1000, 5000 iterations. For our separable approximations, the number of iterations refer to the number of demand samples used. For LSD, SD and CUPPS, the number of iterations refer to the number of cuts used to approximate the recourse function. Having constructed a recourse approximation at iteration k, say $\hat{Q}^k$, we find $(x^k, s^k) \in \arg\max \sum_{i,j\in\mathcal{I}} c_{ij}x_{ij} + \hat{Q}^k(s)$. Then we compute $\sum_{i,j\in\mathcal{I}} c_{ij}x_{ij}^k + \mathbb{E}Q(s^k, \omega)$ in order to measure the performance of the solution (x^k, s^k) provided by the approximation $\hat{Q}^k$.

The results are summarized in table 1. The numbers in the table represent the percent deviation between the optimal objective value and the objective value corresponding to the solution obtained after a certain number of iterations. For all problem instances, we use LSD to find the optimal solution. Table 1 also gives the CPU time per iteration. We present results on ten problems. Six of these problems vary in cost parameters and the last four vary in the dimensionality of the recourse function.

The results indicate that for the problem class we consider, SPAR is able to produce high quality solutions rather quickly and provides consistent performance over different sets of problem parameters. In particular, the consistent performance of SPAR over problems with different numbers of locations may make it appealing for large-scale applications. Nevertheless, our numerical results are limited to a specific problem class and one should be cautious about extending our findings to other problem classes. However, as a result of equation (56), SPAR provides an estimate of the optimality bound at every point it generates.

Considering all of these, SPAR is a promising approach for allocation of indivisible resources under uncertainty, but more comprehensive numerical work is needed before using it in a particular problem context. Finally, we note that due to its simplicity and fast run times, SPAR can be used as an initialization routine for stochastic programming approaches that can exploit high quality feasible solutions. For example, the recourse approximation provided by SPAR can be used to initialize that of SD, or LSD and CUPPS can start by constructing a support of the recourse function at the solution provided by SPAR.

Problem	Method	Number of iterations					Sec./
		25	100	500	1000	5000	iter.
Cost.	SPAR	10.18	0.81	0.21	0.15	0.02	0.07
param.	LSD	29.88	11.68	0.44	0.03		2.41
I	CUPPS	20.74	9.86	5.21	2.39	0	0.37
	SD	47.63	15.64	11.42	7.56	2.48	0.46
Cost.	SPAR	9.28	0.77	0.23	0.26	0.04	0.06
param.	LSD	44.19	8.28	0.49	0.05		2.46
II	CUPPS	49.43	13.58	6.17	1.95	0	0.37
	SD	24.76	17.8	8.56	8.54	1.62	0.46
Cost.	SPAR	10.72	1.64	0.39	0.44	0.1	0.05
param.	LSD	37.07	8.16	0.5	0		2.51
III	CUPPS	36.36	10.99	6.3	2.25	0	0.37
	SD	35.37	17.52	9.14	6.53	2.28	0.45
Cost.	SPAR	10.43	2.41	0.67	0.65	0.08	0.04
param.	LSD	36.18	6.72	0.46	0		2.03
IV	CUPPS	41.79	22.02	9.14	2.94	0	0.64
	SD	57.06	25.4	23.11	12.02	4.45	0.45
Cost.	SPAR	9.58	3.61	0.53	0.74	0.07	0.04
param.	LSD	29.53	11.49	0.33	0		1.68
V	CUPPS	36.28	21.34	9.26	2.5	0	0.64
	SD	25.37	22	23.89	25.93	3.25	0.44
Cost.	SPAR	8.95	4.42	0.75	0.87	0.09	0.04
param.	LSD	40.68	1.91	0			0.93
VI	CUPPS	38.97	6.02	4.77	0.53		0.64
	SD	40.64	11.37	9.22	6.78	1.64	0.43
$\|\mathcal{I} \bigcup \mathcal{C}\|$	SPAR	18.65	7.07	0.48	0.28	0.15	0.00
=10	LSD	3.3	0				0.06
	CUPPS	5.84	0.5	0			0.07
	SD	45.45	11.35	2.3	1.12	0.26	0.10
$\|\mathcal{I} \bigcup \mathcal{C}\|$	SPAR	11.73	2.92	0.34	0.13	0.06	0.02
=25	LSD	19.88	2.14	0			0.26
	CUPPS	8.27	4.33	1.47	0.16		0.31
	SD	40.55	22.22	4.24	4.8	0.95	0.22
$\|\mathcal{I} \bigcup \mathcal{C}\|$	SPAR	9.99	1.18	0.26	0.3	0.05	0.06
=50	LSD	42.56	6.07	0.52	0.04		2.51
	CUPPS	34.93	19.3	5.09	1.38	0	0.37
	SD	43.18	17.94	5.91	6.25	1.02	0.46
$\|\mathcal{I} \bigcup \mathcal{C}\|$	SPAR	8.74	1.2	0.16	0.05	0	0.22
=100*	LSD	74.52	26.21	2.32	0.85	0.02	10.15
	CUPPS	54.59	23.99	14.68	14.13	0.91	1.42
	SD	62.63	40.73	15.22	17.43	9.42	1.30

*Optimal solution not found
Figures represent the deviation from the best objective value known

Table 1: Percent error over optimal with different first stage costs.

References

Birge, J. & Wallace, S. W. (1988), 'A separable piecewise linear upper bound for stochastic linear programs', *SIAM J. Control and Optimization* **26**(3), 1–14. 4, 5

Chen, Z.-L. & Powell, W. (1999), 'A convergent cutting-plane and partial-sampling algorithm for multistage linear programs with recourse', *Journal of Optimization Theory and Applications* **103**(3), 497–524. 4, 26

Cheung, R. K.-M. & Powell, W. B. (2000), 'SHAPE: A stochastic hybrid approximation procedure for two-stage stochastic programs', *Operations Research* **48**(1), 73–79. 4

Cheung, R. & Powell, W. B. (1996), 'An algorithm for multistage dynamic networks with random arc capacities, with an application to dynamic fleet management', *Operations Research* **44**(6), 951–963. 4

Culioli, J.-C. & Cohen, G. (1990), 'Decomposition/coordination algorithms in stochastic optimization', *SIAM Journal of Control and Optimization* **28**, 1372–1403. 4

Ding, X., Puterman, M. L. & Bisi, A. (2002), 'The censored newsvendor and the optimal acquisition of information', *Operations Research* **50**(3), 517–527. 2

Doob (1953), *Stochastic Processes*, John Wiley & Sons. 8, 9

Ermoliev, Y. (1988), Stochastic quasigradient methods, *in* Y. Ermoliev & R. Wets, eds, '*Numerical Techniques for Stochastic Optimization*', Springer-Verlag, Berlin. 4

Frantzeskakis, L. & Powell, W. B. (1990), 'A successive linear approximation procedure for stochastic dynamic vehicle allocation problems', *Transportation Science* **24**(1), 40–57. 4

Godfrey, G. A. & Powell, W. B. (2001), 'An adaptive, distribution-free approximation for the newsvendor problem with censored demands, with applications to inventory and distribution problems', *Management Science* **47**(8), 1101–1112. 1, 5, 20

Higle, J. & Sen, S. (1991), 'Stochastic decomposition: An algorithm for two stage linear programs with recourse', *Mathematics of Operations Research* **16**(3), 650–669. 4, 26

Kleywegt, A. J., Shapiro, A. & Homem-De-Mello, T. (2002), 'The sample average approximation method for stochastic discrete optimization', *SIAM Journal on Optimization* **12**, 479–502. 3

Korf, L. A. & Wets, R. J. B. (2001), 'Random LSC functions: An ergodic theorem', *Mathematics of Operations Research* **26**(2), 421–445. 3

Kushner, H. J. & Yin, G. G. (1997), *Stochastic Approximation Algorithms and Applications*, Springer-Verlag. 4

Laporte, G. & Louveaux, F. (1993), 'The integer L-shaped method for stochastic integer programs with complete recourse', *Operations Research Letters* **13**(3), 133–142. 4

Porteus, E. L. (1990), *Handbooks in Operations Research and Management Science: Stochastic Models*, Vol. 2, North Holland, Amsterdam, chapter Stochastic Inventory Theory, pp. 605–652. 2

Powell, W. B. (1986), 'A stochastic model of the dynamic vehicle allocation problem', *Transportation Science* **20**, 117–129. 4

Powell, W. B. (1987), 'An operational planning model for the dynamic vehicle allocation problem with uncertain demands', *Transportation Research* **21B**, 217–232. 4

Powell, W. B. (1988), A comparative review of alternative algorithms for the dynamic vehicle allocation problem, *in* B. Golden & A. Assad, eds, 'Vehicle Routing: Methods and Studies', North Holland, Amsterdam, pp. 249–292. 4

Powell, W. B. & Carvalho, T. A. (1998), 'Dynamic control of logistics queueing network for large-scale fleet management', *Transportation Science* **32**(2), 90–109. 4

Powell, W. B. & Cheung, R. (1994), 'Stochastic programs over trees with random arc capacities', *Networks* **24**, 161–175. 4

Powell, W. B. & Topaloglu, H. (2003), Stochastic programming in transportation and logistics, *in* A. Ruszczyński & A. Shapiro, eds, '*Handbook in Operations Research and Management Science*, Volume on *Stochastic Programming*', Elsevier, Amsterdam, pp. 555–635. 13

Righter, R. (1989), 'A resource allocation problem in a random environment', *Operations Research* **37**, 329–338. 2

Ruszczyński, A. (1980), 'Feasible direction methods for stochastic programming problems', *Mathematical Programming* **19**, 220–229. 4

Ruszczyński, A. (1986), A method of feasible directions for solving nonsmooth stochastic programming problems, *in* F. Archetti, G. DiPillo & M. Lucertini, eds, 'Stochastic Programming', Springer-Verlag, Berlin, Germany, pp. 258–271. 4

Ruszczyński, A. (1987), 'A linearization method for nonsmooth stochastic programming problems', *Mathematics of Operations Research* **12**(1), 32–49. 4

Ruszczyński, A. (2003), Decomposition methods, *in* A. Ruszczyński & A. Shapiro, eds, '*Handbook in Operations Research and Management Science*, Volume on *Stochastic Programming*', Elsevier, Amsterdam. 4

Shapiro, A. & Homem-De-Mello, T. (2000), 'On the rate of convergence of optimal solutions of Monte Carlo approximations of stochastic programs', *SIAM Journal on Optimization* **11**, 70–86. 3

Topaloglu, H. & Powell, W. B. (2003), 'An algorithm for approximating piecewise linear concave functions from sample gradients', *Operations Research Letters* **31**(1), 66–76. 19

Van Slyke, R. & Wets, R. (1969), 'L-shaped linear programs with applications to optimal control and stochastic programming', *SIAM Journal of Applied Mathematics* **17**(4), 638–663. 4, 26

Wallace, S. (1986), 'Solving stochastic programs with network recourse', *Networks* **16**, 295–317. 4, 5

Wallace, S. (1987), 'A piecewise linear upper bound on the network recourse function', *Mathematical Programming* **38**, 133–146. 4, 5