

On the convergence of stochastic dual dynamic programming and related methods

A. B. Philpott*, Z. Guan†

Abstract

We discuss the almost-sure convergence of a broad class of sampling algorithms for multi-stage stochastic linear programs. We provide a convergence proof based on the finiteness of the set of distinct cut coefficients. This differs from existing published proofs in that it does not require a restrictive assumption.

Keywords: multi-stage stochastic programming, Monte-Carlo sampling, Benders decomposition.

1 Introduction

Multistage stochastic linear programs with recourse are well known in the stochastic programming community, and are becoming more common in applications. The typical approach to solving these problems is to approximate the random variables using a finite set of outcomes forming a scenario tree (see e.g. [1]). This yields a large-scale mathematical programming problem that can be attacked using a suitable algorithm. One algorithm that has been widely applied in energy and logistics settings is the stochastic dual dynamic programming (SDDP) method of Pereira and Pinto [9]. This algorithm constructs feasible dynamic programming policies using an outer approximation of a (convex) future cost function that is computed using Benders cuts. The policies defined by these cuts can be evaluated using simulation, and their performance measured against a lower bound on their expected cost. This

*Corresponding author: Department of Engineering Science, The University of Auckland, Private Bag 92019, Auckland, New Zealand, a.philpott@auckland.ac.nz.

†z.guan@auckland.ac.nz.

provides a convergence criterion that may be applied to terminate the algorithm when the estimated cost of the candidate policy is close enough to its lower bound. The SDDP algorithm has led to a number of related methods (see [2],[3],[4],[6],[10]) that are based on the same essential idea, but seek to improve the method by exploiting the structure of particular applications.

Since its publication in 1991, a number of authors have studied the convergence behaviour of SDDP and related algorithms. In his PhD thesis [3] (and in [4]) Donohue states that “finite convergence of this algorithm follows from the finite convergence of the Nested Decomposition algorithm, since the scenarios from which the optimality cuts are generated are re-sampled at each iteration.” This remark which, strictly speaking, should be a statement of convergence with probability 1, is not accompanied by a formal proof.

The first formal proof of the almost sure convergence of multi-stage sampling algorithms was published by Chen and Powell [2] who derived this for their CUPPS algorithm. This proof was extended by Linowsky and Philpott [8] to cover other multi-stage sampling algorithms (SDDP [9], AND [4], ReSa [6]) that use outer approximation. The convergence proofs in [2] and [8] make use of an unstated assumption regarding the independence of sampled random variables and convergent subsequences of algorithm iterates. The absence of an argument showing that this assumption holds weakens the analysis in these papers, and leaves open the question of convergence in general.

The purpose of this paper is to give a simpler proof of almost sure convergence for a broad class of sampling algorithms that include SDDP, AND, ReSa and CUPPS. Our proof does not require the independence assumption referred to above, but follows the finiteness argument that Donohue alluded to in his thesis, and makes this argument explicit, by basing it on two well-defined properties of the sampling procedure. Although it could be claimed that this proof is already part of the stochastic programming folklore, our contribution is to illuminate the key properties that underly almost sure convergence of these methods.

2 Multistage Benders decomposition

We follow the notation and terminology of [8], and restrict attention to multistage stochastic programs with the following properties.

- (A1) Random quantities appear only on the right-hand side of the linear constraints in each stage.

- (A2) The set Ω_t of random outcomes in each stage $t = 2, 3, \dots, T$ is discrete and finite
 $(\Omega_t = \{\omega_{ti} \mid i = 1, \dots, q_t < \infty\}$ with probabilities $p_{ti} > 0, \forall i$).
- (A3) Random quantities in different stages are independent.
- (A4) The feasible region of the linear program in each stage is non-empty and bounded.

The assumption (A3) includes the case where the right-hand sides of the linear constraints follow an ARMA process, which can be accommodated by augmenting x with variables that model this, along with stagewise independent residuals (see [7]). Under the above assumptions, the multi-stage stochastic linear program can be written in the following form.

Solve the problem defined by

$$\begin{aligned} [LP_1] \quad & Q_1 = \min_{x_1} c_1^\top x_1 + Q_2(x_1) \\ \text{subject to} \quad & A_1 x_1 = b_1, \\ & x_1 \geq 0, \end{aligned}$$

where for all $t = 2, \dots, T$,

$$Q_t(x_{t-1}) = \sum_{i=1}^{q_t} p_{ti} Q_t(x_{t-1}, \omega_{ti}),$$

$Q_t(x_{t-1}, \omega_{ti})$ is defined by the problem

$$\begin{aligned} [LP_t] \quad & Q_t(x_{t-1}, \omega_t) = \min_{x_t} c_t^\top x_t + Q_{t+1}(x_t) \\ \text{subject to} \quad & A_t x_t = \omega_t - B_{t-1} x_{t-1}, \\ & x_t \geq 0, \end{aligned}$$

and we set $Q_{T+1} \equiv 0$.

The problem $[LP_t]$ depends on the choice of ω_t and x_{t-1} , and so we could write $[LP_t(x_{t-1}, \omega_t)]$, though we choose to suppress this dependence in the notation. By Assumption (A3), $[LP_t]$ is independent of $\omega_{t-1}, \omega_{t-2}, \dots$.

In the class of sampling algorithms that we consider in this paper the functions $Q_t(x_{t-1})$ in each stage are approximated by the maximum of a collection of linear functions, each of which is called a *cut*. In each iteration $k = 1, 2, \dots$, the type of algorithm we consider computes a set of feasible solutions $\{x_t^k : t = 1, \dots, T-1\}$, and a set of cuts, one for each stage $t = 1, \dots, T-1$. This gives rise to a sequence of approximate problems $[AP_t^k]$, $k = 1, 2, \dots$, for each stage. These are defined as follows:

For $t = 1$, we solve the linear program

$$\begin{aligned} [AP_1^k] \quad & C_1^k = \min_{x_1, \theta_2} c_1^\top x_1 + \theta_2 \\ \text{subject to} \quad & A_1 x_1 = b_1, \\ & \theta_2 + (\beta_2^j)^\top x_1 \geq \alpha_{2,j}, \quad j = 0, \dots, k-1, \\ & x_1 \geq 0, \end{aligned}$$

and, for $t = 2, \dots, T-1$, we solve

$$\begin{aligned} [AP_t^k] \quad & C_t^k(x_{t-1}^k, \omega_t) = \min_{x_t, \theta_{t+1}} c_t^\top x_t + \theta_{t+1} \\ \text{subject to} \quad & A_t x_t = \omega_t - B_{t-1} x_{t-1}^k, \\ & \theta_{t+1} + (\beta_{t+1}^j)^\top x_t \geq \alpha_{t+1,j}, \quad j = 0, \dots, k-1, \\ & x_t \geq 0. \end{aligned}$$

Finally for every k , we set $[AP_T^k] = [LP_T]$. The problems $[AP_t^k]$ are approximations of $[LP_t]$ in the sense that $Q_{t+1}(x_t)$ is approximated (below) by the polyhedral function

$$\max_{j=0, \dots, k-1} \{ \alpha_{t+1,j} - (\beta_{t+1}^j)^\top x_t \}.$$

This means that any solution to $[AP_t^k]$ has a value that is a lower bound on the optimal value of $[LP_t]$.

For all stages, the first cut ($j = 0$) is set as the trivial cut $\theta_{t+1} \geq -\infty$. We use the notation $C_t^k(x_{t-1})$ to denote $\sum_{i=1}^{q_t} p_{ti} C_t^k(x_{t-1}, \omega_t)$. In the last stage, T , we have $[AP_T^k] = [LP_T]$, and so for every x_{T-1} and ω_T

$$C_T^k(x_{T-1}, \omega_T) = Q_T(x_{T-1}, \omega_T), \quad k = 1, 2, \dots$$

Since cuts are added from one iteration to the next, and no cuts are taken out, the optimal values of $[AP_t^k]$ form a monotonic sequence, i.e. for $k = 1, 2, \dots$

$$C_t^{k+1}(x_{t-1}, \omega_t) \geq C_t^k(x_{t-1}, \omega_t), \quad t = 2, 3, \dots, T,$$

and

$$C_1^{k+1} \geq C_1^k.$$

Observe that under Assumption (A4),

$$\{x_t \mid A_t x_t = \omega_t - B_{t-1} x_{t-1}^k, \quad x_t \geq 0\}$$

is nonempty and bounded so $[AP_t^k]$ always has a nonempty feasible set (with θ_{t+1} chosen large enough) and hence an optimal solution. Thus its dual has

an optimal solution (π_t, ρ_t) , where π_t corresponds to the equality constraints, and ρ_t corresponds to the cut constraints. Furthermore, by Assumption (A1), the set of extreme points of the dual of $[AP_t^k]$ is independent of the outcomes of the random quantities, which allows us to construct a valid cut at each stage based on an assembled collection \mathcal{D}_t^k of extreme-point dual solutions from different samples.

Initially at iteration $k = 0$, $\mathcal{D}_t^0 = \emptyset$. At any subsequent iteration k the coefficients of the cuts at each stage $t = 1, 2, \dots, T - 1$, are calculated as follows.

Cut Calculation Algorithm (CCA)

1. Choose a sample $\Omega_t^k \subseteq \Omega_t$, solve $[AP_t^k]$ for all $\omega_{ti} \in \Omega_t^k$, and add the optimal extreme-point dual solutions to \mathcal{D}_t^k .
2. Let $(\pi_t^i(x_{t-1}^k), \rho_t^i(x_{t-1}^k))$ be the best dual solution in \mathcal{D}_t^k for $[AP_t^k]$ for each $\omega_{ti} \in \Omega_t$, that is, if $t < T$ then

$$\begin{aligned} & (\pi_t^i(x_{t-1}^k), \rho_t^i(x_{t-1}^k)) \\ = & \arg \max \{ \pi_t^\top (\omega_{ti} - B_{t-1} x_{t-1}^k) + \rho_t^\top \alpha_{t+1}^{k-1} \mid (\pi_t, \rho_t) \in \mathcal{D}_t^k \}, \end{aligned}$$

and

$$\pi_T^i(x_{T-1}^k) = \arg \max \{ \pi_T^\top (\omega_{Ti} - B_{T-1} x_{T-1}^k) \mid \pi_T \in \mathcal{D}_T^k \}$$

otherwise.

3. The cut has the formula

$$\theta_t \geq \alpha_{t,k} - (\beta_t^k)^\top x_{t-1}$$

where

$$\begin{aligned} \beta_t^k &= \sum_{i=1}^{q_t} p_{ti} B_{t-1}^\top \pi_t^i(x_{t-1}^k) && \text{for } 2 \leq t \leq T, \\ \alpha_{t,k} &= \sum_{i=1}^{q_t} p_{ti} [\omega_{ti}^\top \pi_t^i(x_{t-1}^k) + (\alpha_{t+1}^{k-1})^\top \rho_t^i(x_{t-1}^k)] && \text{for } 2 \leq t \leq T - 1, \\ \alpha_{T,k} &= \sum_{i=1}^{q_T} p_{Ti} \omega_{Ti}^\top \pi_T^i(x_{T-1}^k). \end{aligned}$$

CCA is the same algorithm as that used to define a *sampled cut* in [8]. Observe that $\alpha_{t,k}$ is a scalar, whereas α_{t+1}^{k-1} denotes a $(k-1)$ -dimensional vector. This means that the dimensions of α_{t+1}^{k-1} and $\rho_t^i(x_{t-1}^k)$ are increasing as the iteration count k increases, and thus the collection of extreme-point solutions of the dual of $[AP_t^k]$ may be infinite. On the other hand, the collection of distinct values of $(\beta_t^k, \alpha_{t,k})$ is provably finite, as we show in the following lemma.

Lemma 1 *For each $t = 2, 3, \dots, T$, define the set*

$$\mathcal{G}_t^k = \{(\beta_t^j, \alpha_{t,j}) : j = 1, 2, \dots, k-1\}.$$

Then for any sequence \mathcal{G}_t^k , $k = 1, 2, \dots$ generated by the repeated application of CCA there exists m_t such that for all k

$$|\mathcal{G}_t^k| \leq m_t.$$

Furthermore, there exists k_t , so that if $k > k_t$ then $\mathcal{G}_t^k = \mathcal{G}_t^{k_t}$.

Proof. Consider any realization of the sequence \mathcal{G}_t^k , $k = 1, 2, \dots$ generated by the repeated application of CCA. We use induction on t to construct m_t such that $|\mathcal{G}_t^k| \leq m_t$. The second part of the lemma follows immediately. First at T , $\rho_T = 0$ and π_T is an extreme point of $\{\pi \mid A_T^\top \pi \leq c_T\}$ of which there are at most m_{T+1} , say. Then the cut coefficients

$$\begin{aligned} \alpha_{T,k} &= \sum_{i=1}^{q_T} p_{Ti} \omega_{Ti}^\top \pi_T^i(x_{T-1}^k), \\ \beta_T^k &= \sum_{i=1}^{q_T} p_{Ti} B_{T-1}^\top \pi_T^i(x_{T-1}^k), \end{aligned}$$

each can only take at most $m_{T+1}^{q_T}$ values, and thus if $m_T = m_{T+1}^{2q_T}$, then for all k

$$|\mathcal{G}_T^k| \leq m_T.$$

Now suppose at t that there exists m_{t+1} such that for all k

$$|\mathcal{G}_{t+1}^k| \leq m_{t+1}.$$

It follows that there exists k_{t+1} , so that if $k > k_{t+1}$ then $\mathcal{G}_{t+1}^k = \mathcal{G}_{t+1}^{k_{t+1}}$ and the cut at iteration $k > k_{t+1}$ is a repeat of some cut in the existing cuts. Consider the feasible region of the dual of $[AP_t^k]$, namely

$$\mathcal{H}_t^k = \{(\pi_t, \rho_t) \mid A_t^\top \pi_t + \sum_{j=1}^{k-1} \beta_{t+1}^j \rho_t^j \leq c_t, \quad \sum_{j=1}^{k-1} \rho_t^j = 1, \quad \rho_t \geq 0\}.$$

If $k > k_{t+1}$ then any extreme point (π_t^k, ρ_t^k) of \mathcal{H}_t^k corresponds to an extreme point (π, ρ) of $\mathcal{H}_t^{k_{t+1}}$ with the same dual objective value, obtained by choosing $\pi = \pi_t^k$ and basic columns β_{t+1}^j for $j < k_{t+1}$ that match the basic columns β_{t+1}^j , $k_{t+1} \leq j < k$. This is because each latter column β_{t+1}^j and its cost coefficient $\alpha_{t+1,j}$ is a duplicate of some $(\beta, \alpha) \in \mathcal{G}_{t+1}^{k_{t+1}}$. Since there are a finite number, say e_t , of extreme point solutions to $\mathcal{H}_t^{k_{t+1}}$, there are at most e_t distinct values of

$$[\omega_{ti}^\top \pi_t^i(x_{t-1}^k) + (\alpha_{t+1}^{k-1})^\top \rho_t^i(x_{t-1}^k)]$$

and so $(e_t)^{q_t}$ distinct values of

$$\alpha_{t,k} = \sum_{i=1}^{q_t} p_{ti} [\omega_{ti}^\top \pi_t^i(x_{t-1}^k) + (\alpha_{t+1}^{k-1})^\top \rho_t^i(x_{t-1}^k)],$$

Similarly,

$$\beta_t^k = \sum_{i=1}^{q_t} p_{ti} B_{t-1}^\top \pi_t^i(x_{t-1}^k),$$

can take at most $(e_t)^{q_t}$ values and so if $m_t = (e_t)^{2q_t}$ then

$$|\mathcal{G}_t^k| \leq m_t,$$

which proves the result. ■

Linowsky and Philpott [8] define a class of sampling-based decomposition algorithms, the Multi-stage Sampled Benders Decomposition (MSBD), which includes SDDP, AND, ReSa and CUPPS. Here we define a different class of algorithms that we call Dynamic Outer Approximation Sampling Algorithms (DOASA). For DOASA the sampling procedures must satisfy two distinct properties FPSP and BPSP that are defined below. We make use of the terminology *scenario* to denote an element of $\prod_{t=2}^{T-1} \Omega_t$ indexed by j so

$$\prod_{t=2}^{T-1} \Omega_t = \{\omega(j) \mid j = 1, 2, \dots, \prod_{t=2}^{T-1} q_t\}.$$

Algorithms in the DOASA class perform the following steps:

Step 0: (Initialization) Set $k = 1$.

Step 1: (Forward pass)

Sample a single outcome ω_t of the random variable in each stage $t = 2, 3, \dots, T-1$, to give a single scenario $\{\omega_t^k\}$ satisfying FPSP. For each stage $t = 1, 2, \dots, T-1$, compute the primal solution (x_t^k, θ_{t+1}^k) of the problem $[AP_t^k]$.

Step 2: (Cut Generation)

For each stage $t = T, T - 1, \dots, 2$, apply CCA to generate a cut at x_{t-1}^k with sample Ω_t^k satisfying BPSP.

Step 3: Set $k = k + 1$ and go to Step 1.

We now formally state the sampling properties as follows:

Forward Pass Sampling Property (FPSP):

For each $j = 1, 2, \dots, \prod_{t=2}^{T-1} q_t$, with probability 1

$$|\{k : \{\omega_t^k \mid t = 2, 3, \dots, T - 1\} = \omega(j)\}| = \infty.$$

Backward Pass Sampling Property (BPSP):

For each $t = 2, 3, \dots, T$ and $i = 1, 2, \dots, q_t$, with probability 1

$$|\{k : \omega_{ti} \in \Omega_t^k\}| = \infty.$$

FPSP states that each scenario $\omega(j)$ is traversed infinitely many times with probability 1 in the forward pass. BPSP states that each scenario outcome ω_{ti} is visited infinitely many times with probability 1 in the backward pass. There are many sampling methods satisfying these two properties. For example, consider independently sampling a single outcome in each stage with a positive probability for each ω_{ti} in the forward pass and backward pass respectively. Then by the Borel-Cantelli lemma (see [5]) this method satisfies both properties. Another sampling method that satisfies FPSP and BPSP is to repeat an exhaustive enumeration of each scenario $\omega(j)$, $j = 1, 2, \dots, \prod_{t=2}^{T-1} q_t$ in both the forward pass and the backward pass, although such a method would be prohibitively expensive in all but the smallest examples.

3 Convergence of DOASA algorithms

Previous published results in [2] and [8] give proofs for the almost sure convergence of CUPPS and MSBD respectively. The proofs in both of these papers require an important but unstated assumption. Here we state this assumption formally and discuss it.

Let the iterations of the algorithm be indexed by $\mathcal{N} = \{1, 2, \dots\}$ and suppose $t \in \{1, \dots, T - 1\}$. Let $\{\omega_t^n, x_t^n\}_{n \in \mathcal{N}}$ be the sequence generated by the sampling algorithm at stage t .

Assumption 1: For any infinite subsequence $\{x_t^k\}_{k \in \mathcal{K}}$ of $\{x_t^n\}_{n \in \mathcal{N}}$ there exists a convergent subsequence $\{x_t^j\}_{j \in \mathcal{J}}$ that is independent of $\{\omega_{t+1}^j\}_{j \in \mathcal{J}}$.

Remark 4.1 in [2] correctly claims that if \mathcal{N} is infinite then with probability one \mathcal{N} has an infinite subset \mathcal{N}_{ti} corresponding to draws of outcome ω_{ti} for any $i = 1, \dots, q_t$ and $t = 2, \dots, T$. This follows by an application of the Borel-Cantelli lemma, because each ω_t^n in $\{\omega_t^n\}_{n \in \mathcal{N}}$ is independently sampled and $\Pr[\omega_t^n = \omega_{ti}] > 0$.

However, the situation becomes more subtle in the proof of Lemma 5.2 in [2]. Here the authors claim that for any infinite subset \mathcal{K} of \mathcal{N} , there exists an infinite subset \mathcal{J} with a convergent subsequence $\{x_{T-1}^j\}_{j \in \mathcal{J}}$ such that with probability one there exists an infinite subset \mathcal{J}_i of \mathcal{J} corresponding to draws of each sample ω_{Ti} for $i = 1, \dots, q_T$. The convergent subsequence $\{x_{T-1}^j\}_{j \in \mathcal{J}}$ in this lemma is constructed using the assumed compactness of the set X in which x_{T-1} lies. Of course, compactness guarantees a convergent subsequence $\{x_{T-1}^j\}_{j \in \mathcal{J}}$ of $\{x_{T-1}^k\}_{k \in \mathcal{K}}$, but it cannot be deduced from this and Remark 4.1 in [2] that there are an infinite number of ω_{Ti} in $\{\omega_T^j\}_{j \in \mathcal{J}}$ for every $i = 1, \dots, q_T$. (The problem here is that for every *convergent* subsequence it might be the case that there are only finitely many ω_{Ti} for some $i = 1, \dots, q_T$, and this possibility needs to be ruled out somehow.)

In claiming the independence of the sampling procedure from the convergence of the subsequence, the authors of [2] are making an implicit assumption (Assumption 1), which is needed to make the proof of Lemma 5.2 valid. The proof in [8] is based on Lemma 5.2 in [2], and so it is also flawed in the absence of Assumption 1.

We now proceed to give a direct proof of almost sure convergence that does not rely on Assumption 1. The new proof formalizes the assertion by Donohue [3] that convergence follows from re-sampling. To understand how the proof works consider an approach (DOASA- N) in which a finite set of N scenarios is sampled in advance, and the forward pass of the algorithm traverses each of these scenarios.

DOASA- N

Step 0: (Initialization) Set $k = 1$. For $s = 1$ to N , select at each stage $t = 2, 3, \dots, T - 1$, a single outcome ω_{st} of the random variable to give a set of N scenarios.

Step 1: (Forward pass)

For each scenario s , and stage $t = 1, 2, \dots, T - 1$, compute the primal solution $(x_{st}^k, \theta_{s,t+1}^k)$ of the problem $[AP_t^k]$.

Step 2: (Cut Generation)

For each stage $t = T, T-1, \dots, 2$, apply CCA to generate N cuts at the states $x_{s,t-1}^k$ with samples $\Omega_{s,t}^k$, $s = 1, 2, \dots, N$.

Step 3: Set $k = k + 1$ and go to Step 1.

Lemma 2 *In every realization of iterations, DOASA- N converges in a finite number of iterations to a policy giving value $\lim_k C_1^k$ which is at most equal to the optimal expected cost of $[LP_1]$.*

Proof. For each $s = 1, 2, \dots, N$, since $k = 1, 2, \dots$ and one cut is constructed in each iteration k , then by Lemma 1, for $t \in \{2, \dots, T\}$, there exists $k_{s,t}$, so that if $k > k_{s,t}$ then $\mathcal{G}_t^k = \mathcal{G}_t^{k_{s,t}}$ and thus there is no further change in the cuts defining $\mathcal{C}_{s,t}^k(x_{s,t-1})$, that is, for every $x_{s,t-1}$

$$\max_{j=0, \dots, k-1} \{\alpha_{t,j} - (\beta_t^j)^\top x_{s,t-1}\} = \max_{j=0, \dots, k_{s,t}-1} \{\alpha_{t,j} - (\beta_t^j)^\top x_{s,t-1}\}.$$

For $t \in \{2, \dots, T\}$, if we choose $k_t = \max_{s=1}^N \{k_{s,t}\}$, then for each $k > k_t$ there is no change in the cuts defining $\mathcal{C}_t^k(x_{t-1})$, that is, for every x_{t-1}

$$\max_{j=0, \dots, k-1} \{\alpha_{t,j} - (\beta_t^j)^\top x_{t-1}\} = \max_{j=0, \dots, k_t-1} \{\alpha_{t,j} - (\beta_t^j)^\top x_{t-1}\}.$$

Thus all solutions (x_1^k, θ_2^k) to $[AP_1^k]$ are the same for $k > k_2$, as are all solutions (x_t^k, θ_{t+1}^k) to $[AP_t^k]$, $t = 2, 3, \dots, T$, so DOASA- N terminates after iteration k_2 .

It is easy to see that for every k the optimal value of $[AP_1^k]$ is a lower bound on the optimal expected cost of $[LP_1]$. ■

A special case of DOASA- N uses the universe of $N = \prod_{t=2}^{T-1} q_t$ scenarios.

Lemma 3 *Under BPSP, DOASA- N with the universe of scenarios converges with probability 1 to an optimal solution to $[LP_1]$ in a finite number of iterations.*

Proof. From Lemma 2 in every realization of iterations DOASA- N will converge in a finite number of steps to a policy that has $\lim_k C_1^k$ giving a lower bound on the true expected cost. Now consider a realization of DOASA- N iterations, and denote the limiting policy by $(\bar{x}_1, \bar{x}_2(\omega_2), \bar{x}_3(\omega_2, \omega_3), \dots)$, which is obtained at iteration \bar{k} , say. For any scenario $\omega_2, \omega_3, \dots, \omega_T$, we denote $\bar{x}_t(\omega_2, \dots, \omega_t)$ by $\bar{x}_t(\omega)$. We claim that for every $k > \bar{k}$, and any scenario ω ,

$$C_T^k(\bar{x}_{T-1}(\omega)) = \mathcal{Q}_T(\bar{x}_{T-1}(\omega)), \quad (1)$$

with probability 1, which implies $C_T^k(\bar{x}_{T-1}(\omega), \omega_T) = Q_T(\bar{x}_{T-1}(\omega), \omega_T)$ for all ω_T . Otherwise for some particular outcome $\hat{\omega}_T$, we have $\hat{\omega}_T \notin \Omega_t^k$, for every $k > \bar{k}$, with positive probability which violates BPSP.

Now we claim that if $k > \bar{k}$ then for every scenario ω

$$\mathcal{C}_{T-1}^k(\bar{x}_{T-2}(\omega)) = \mathcal{Q}_{T-1}(\bar{x}_{T-2}(\omega)). \quad (2)$$

Otherwise for some particular outcome $\hat{\omega}_{T-1}$,

$$\mathcal{C}_{T-1}^k(\bar{x}_{T-2}(\omega), \hat{\omega}_{T-1}) < Q_{T-1}(\bar{x}_{T-2}(\omega), \hat{\omega}_{T-1}). \quad (3)$$

But

$$\begin{aligned} \mathcal{C}_{T-1}^k(\bar{x}_{T-2}(\omega), \hat{\omega}_{T-1}) = & \min_{x_{T-1}, \theta_T} c_{T-1}^\top x_{T-1} + \theta_T \\ \text{subject to} & A_{T-1} x_{T-1} = \hat{\omega}_{T-1} - B_{T-2} \bar{x}_{T-2}(\omega), \\ & \theta_T + (\beta_T^j)^\top x_{T-1} \geq \alpha_{T,j}, \quad j = 0, \dots, k-1, \\ & x_{T-1} \geq 0, \end{aligned}$$

which has optimal solution

$$(x_{T-1}^*, \theta_T^*) = (\bar{x}_{T-1}(\omega), \max_{j=0, \dots, k-1} \{\alpha_{T,j} - (\beta_T^j)^\top \bar{x}_{T-1}(\omega)\})$$

with $\omega_{T-1} = \hat{\omega}_{T-1}$.

If $\theta_T^* < \mathcal{C}_T^k(x_{T-1}^*)$, then for any $k > \bar{k}$

$$\max_{j=0, \dots, k-1} \{\alpha_{T,j} - (\beta_T^j)^\top \bar{x}_{T-1}(\omega)\} < \mathcal{C}_T^k(x_{T-1}^*) = \mathcal{Q}_T(\bar{x}_{T-1}(\omega)) \quad (4)$$

by (1). But by BPSP we have with probability 1 that for each ω_T there is some $k(\omega_T) > \bar{k}$ with $\omega_T \in \Omega_T^{k(\omega_T)}$. If we let \hat{k} denote the maximum of the $k(\omega_T)$ then the height of the cut at $\bar{x}_{T-1}(\omega)$ evaluated at iteration \hat{k} is $\mathcal{Q}_T(\bar{x}_{T-1}(\omega))$ contradicting (4) (see Figure 1). Thus we have

$$\theta_T^* = \mathcal{C}_T^k(x_{T-1}^*) = \mathcal{Q}_T(x_{T-1}^*)$$

and

$$\mathcal{C}_{T-1}^k(\bar{x}_{T-2}(\omega), \hat{\omega}_{T-1}) = c_{T-1}^\top x_{T-1}^* + \mathcal{Q}_T(x_{T-1}^*) = Q_{T-1}(\bar{x}_{T-2}(\omega), \hat{\omega}_{T-1})$$

contradicting (3), thereby demonstrating (2). Observe that since $\hat{\omega}_{T-1}$ was arbitrary this shows that $\bar{x}_{T-1}(\omega)$ solves $[LP_{T-1}(\bar{x}_{T-2}(\omega), \omega_{T-1})]$ for any ω_{T-1} .

In a similar way, it is easy to show by induction that $\bar{x}_{t-1}(\omega)$ solves $[LP_{t-1}(\bar{x}_{t-2}(\omega), \omega_{t-1})]$ thus demonstrating that $(\bar{x}_1, \bar{x}_2(\omega_2), \bar{x}_3(\omega_2, \omega_3), \dots)$ is an optimal policy. ■

We now return to the DOASA class of algorithms, in which a single scenario is re-sampled in each forward pass, in contrast to the methods above when these are sampled once and then fixed.

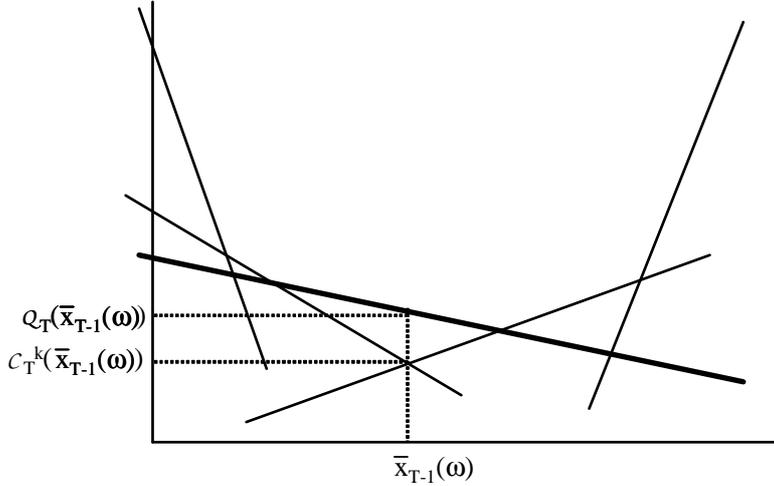


Figure 1: A new cut shown in bold would be created if $\theta_T^* < C_T^k(x_{T-1}^*)$.

Theorem 4 *Under FPSP and BPSP, DOASA converges with probability 1 to an optimal solution to $[LP_1]$ in a finite number of iterations.*

Proof. By FPSP, each scenario in the finite collection of $N = \prod_{t=2}^{T-1} q_t$ scenarios will occur an infinite number of times in the course of the algorithm with probability 1. Thus with probability 1, DOASA will contain a sequence of iterations that are equivalent to DOASA- N applied to the universe of scenarios. We may then apply Lemma 3 which shows that with probability 1, DOASA will converge in a finite number of steps to an optimal solution to $[LP_1]$ in a finite number of iterations. ■

4 Discussion

The proof of convergence in [8] makes some different assumptions from FPSP and BPSP, namely the Cut Sampling Property and the Sample Intersection Property. The Cut Sampling Property (CSP) states that there are only a finite number of iterations in the algorithm where Ω_t^k is empty. Since we are investigating convergence as $k \rightarrow \infty$, CSP is effectively the same as assuming that Ω_t^k is nonempty for all k . The Sample Intersection Property (SIP) states that for any t , each $\omega_{ti} \in \Omega_t$ and each k (given $\Omega_t^k \neq \emptyset$),

$$\Pr[(\omega_{ti} \in \Omega_t^k) \cap (\omega_t^k = \omega_{ti})] > 0.$$

The following Lemma shows that SIP is sufficient to guarantee FPSP and BPSP if it is accompanied by independent sampling in the forward pass and the backward pass.

Lemma 5 *Given independent sampling in the forward pass, SIP implies FPSP. Given independent sampling in the backward pass, SIP implies BPSP.*

Proof. By SIP, for each $\omega_{ti} \in \Omega_t$ and each k (given $\Omega_t^k \neq \emptyset$),

$$\Pr[\omega_t^k = \omega_{ti}] > 0, \quad (5)$$

$$\Pr[\omega_{ti} \in \Omega_t^k] > 0. \quad (6)$$

By (5) and independent sampling in the forward pass, for any scenario $\omega(j)$ with $\omega_{ti} \in \omega(j)$, $t = 2, 3, \dots, T-1$,

$$\Pr[\{\omega_t^k\} = \omega(j)] = \prod_{t=2}^{T-1} \Pr[\omega_t^k = \omega_{ti}] > 0.$$

Then with independent sampling in the forward pass, by the Borel-Cantelli lemma, there are infinite traversals of each scenario $\omega(j)$, $j = 1, 2, \dots, \prod_{t=2}^{T-1} q_t$ with probability 1, and thus FPSP is satisfied. Given (6) and independent sampling in the backward pass, by the Borel-Cantelli lemma, there are infinite visits to each ω_{ti} with probability 1, and thus BPSP is satisfied. ■

The CUPPS algorithm in [2] uses independent sampling in the forward pass, and cuts computed using $\Omega_t^k = \{\omega_t^k\}$. This is easily seen to satisfy SIP, and FPSP and BPSP. Observe, however, that SIP is not necessary for FPSP and BPSP to hold. A version of CUPPS in which cuts are computed using $\Omega_t^k = \Omega_t \setminus \{\omega_t^k\}$ does not satisfy SIP, but it does satisfy FPSP and BPSP. The algorithms SDDP, AND, and ReSa all use independent sampling in the forward pass, and set $\Omega_t^k = \Omega_t$. (We are assuming here that SDDP re-samples in its forward pass. Some commercial implementations of SDDP do not re-sample and so are more akin to DOASA-N than DOASA.) Thus BPSP is trivially true, and FPSP follows by the Borel-Cantelli lemma. These algorithms also satisfy SIP trivially.

Finally, Lemma 8 in [8] asserts that CSP, SIP and independent sampling in the forward pass are sufficient for almost sure convergence. As discussed above there is an implicit independence assumption in the proof of Lemma 8. It is tempting to suppose that independent sampling in the forward pass and SIP give BPSP, which would make Lemma 8 true. However this is not true in general. Thus, in the absence of independent sampling in the backward pass, Lemma 8 in [8] remains unproven.

Acknowledgements: The work of the first author was supported by FRST grant UOA0203. The authors would like to acknowledge the helpful comments of Geoffrey Pritchard and Claudia Sagastizábal.

References

- [1] J.R. Birge, F. Louveaux, Introduction to Stochastic Programming, Springer Verlag, New York, 1997.
- [2] Z.L. Chen, W.B. Powell, Convergent Cutting Plane and Partial-Sampling Algorithm for Multistage Stochastic Linear Programs with Recourse, *Journal of Optimization Theory and Applications* 102 (1999) 497-524.
- [3] C.J. Donohue, Stochastic Network Programming and the Dynamic Vehicle Allocation Problem, PhD Dissertation, University of Michigan, Ann Arbor, Michigan, 1996.
- [4] C.J. Donohue, J.R. Birge, The Abridged nested Decomposition Method for Multistage Stochastic Linear Programs with Relatively Complete Recourse, *Algorithmic Operations Research* 1 (2006) 20-30.
- [5] G.R. Grimmett, D.R. Stirzaker, Probability and Random Processes, Oxford University Press, Oxford, UK, 1992.
- [6] M. Hindsberger, A.B. Philpott, ReSa: A Method for Solving Multistage Stochastic Linear Programs, SPIX Stochastic Programming Symposium, Berlin, 2001.
- [7] G. Infanger, D.P. Morton, Cut Sharing for Multistage Stochastic Linear Programs with Interstage Dependency, *Math. Programming* 75 (2006) 241-256.
- [8] K. Linowsky, A.B. Philpott, On the Convergence of Sampling-Based Decomposition Algorithms for Multistage Stochastic Programs, *Journal of Optimization Theory and Applications* 125 (2005) 349-366.
- [9] M.V.F. Pereira, L.M.V.G. Pinto, Multi-Stage Stochastic Optimization Applied to Energy Planning, *Math. Programming* 52 (1991) 359-375.
- [10] W. Powell, A. Ruszczyński, H. Topaloglu, Learning Algorithms for Separable Approximations of Discrete Stochastic Optimization Problems, *Math. Oper. Res.* 29 (2004) 814-836.