MATHEMATICS OF OPERATIONS RESEARCH
Vol. 00, No. 0, Xxxxx 0000, pp. 000–000
ISSN 0364-765X | EISSN 1526-5471 | 00 | 0000 | 0001

INFORMS
DOI 10.1287/xxxx.0000.0000

On the convergence of decomposition methods for multi-stage stochastic convex programs

P. Girardeau

P. Girardeau, University of Auckland, Dept. of Engineering Science, 70 Symonds Street, Auckland 1021, New Zealand. Also with EDF R&D, Clamart, France, pierre.girardeau@ensta.org

A. B. Philpott

A. B. Philpott, University of Auckland, Dept. of Engineering Science, 70 Symonds Street, Auckland 1021, New Zealand, a.philpott@auckland.ac.nz

We prove the almost-sure convergence of a class of sampling-based nested decomposition algorithms for multi-stage stochastic convex programs in which the stage costs are piecewise smooth functions of the decisions, and uncertainty is modelled by a scenario tree. As a special case, our results imply the almost-sure convergence of SDDP, CUPPS and DOASA when applied to problems with general convex cost functions.

Key words: Stochastic Programming; Dynamic Programming; Stochastic Dual Dynamic Programming algorithm; Monte-Carlo sampling; Benders decomposition
MSC2000 subject classification: Primary: 90C14; secondary: 90C39

1. Introduction

Multistage stochastic programs with recourse are well known in the stochastic programming community, and are becoming more common in applications. We are motivated in this paper by applications in which the stage costs are piecewise smooth convex functions of the decisions. Production functions are often modelled as smooth concave functions of allocated resources. For example Finardi and da Silva [4] use this approach to model hydro electricity production as a concave function of water flow. Smooth value functions also arise when one maximizes profit with linear demand functions (see e.g. [10]) giving a concave quadratic objective or when smooth coherent risk measures are used to model aversion to risk in multistage problems [12].

Having general convex stage costs does not preclude the use of cutting plane algorithms for attacking these problems. The most well-known of these is the stochastic dual dynamic programming (SDDP) algorithm of Pereira and Pinto [8]. 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 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 [1, 2, 3, 6, 9] that are based on the same essential idea, but seek to improve the method by exploiting the structure of particular applications. We call these methods DOASA for (dynamic outer-approximation sampling algorithms) but they are now generically named SDDP methods.

SDDP methods are known to converge almost surely on a finite scenario tree when the stage problems are linear programs. The first formal proof of such a result was published by Chen and Powell [1] who derived this for their CUPPS algorithm. This proof was extended by Linowsky and Philpott [7] to cover other SDDP algorithms. The convergence proofs in [1] and [7] make use of an unstated assumption regarding the independence of sampled random variables and convergent subsequences of algorithm iterates. This assumption was identified by Philpott and Guan [9], who gave a simpler proof of almost sure convergence of SDDP methods based on the finite convergence of the nested decomposition algorithm (see [2]). This does not require the unstated assumption, but exploits the fact that the collection of subproblems to be solved has a finite number of dual extreme points. This begs the question of whether SDDP methods will converge almost surely for general convex stage problems, where the value functions may admit an infinite number of subgradients.

In this paper we propose a different approach from the one in [1] and [7] and show how a proof of convergence for general sampling-based nested decomposition algorithms on finite scenario trees can be established for models with convex subproblems (which may not have polyhedral value functions). As a special case, our approach proves that SDDP methods for multi-stage stochastic convex programming will converge almost surely as long as the sampling process visits every node of the tree an infinite number of times with an appropriate sampling procedure. Our result also covers other sampling procedures that, to the best of our knowledge, have not been applied to solution procedures for multistage stochastic programs.

The result we prove works in the space of state variables expressed as random variables adapted to the filtration defined by the scenario tree. Because this tree has a finite number of nodes, this space is compact, and so we may extract convergent subsequences for any infinite sequence of states. Unlike the arguments in [1] and [7], these subsequences are not explicitly constructed, and so we can escape the need to assume properties of them that we wish to be inherited from independent sampling.

The paper is laid out as follows. We first consider a deterministic multi-stage problem, in which the proof is easily understandable. This is then extended to a stochastic problem in §3. We close with some remarks about the convergence of sampling algorithms.

2. The deterministic case We begin with a lemma that shows that Kelley's cutting plane method converges when applied to the optimization problem:

$$W^* := \min_{u \in \mathcal{U}} W(u),$$

where $\mathcal{U}$ is a convex subset of $\mathbb{R}^m$, and W is a convex function.

Even though the following lemma is not directly used in the more complex results that follow, the main ideas on which the proofs rely will be the same, but presented here in a simpler framework. We believe the reader will find it convenient to already have the scheme of the proof in mind when attacking the more important results later on.

Kelley's method generates a sequence of iterates $(u^j)_{j \in \mathbb{N}}$ by solving, at each iteration, a piecewise linear model of the original problem. The model is then enhanced by adding a cutting plane based on the value $W(u^j)$ and subgradient g^j of W at u^j. The model at iteration k is denoted by

$$W^k(u) := \max_{1 \le j \le k} \left(W\left(u^j\right) + \left\langle g^j, u - u^j \right\rangle \right),$$

and $\theta^k := \min_{u \in \mathcal{U}} W^k(u) = W^k(u^{k+1})$. We have the following result.

Lemma 1. *If W has bounded subgradients and $\mathcal{U}$ is compact then*

$$\lim_{k \to +\infty} W\left(u^k\right) = W^*.$$

T his proof is taken from Ruszczynski [11, Theorem 7.7]. Let K_ε be the set of indices k such that $W^* + \varepsilon < W(u^k) < +\infty$. The proof consists in showing that K_ε is finite.

Suppose $k_1, k_2 \in K_\varepsilon$ and k_1 is strictly smaller than k_2. We have that $W(u^{k_1}) > W^* + \varepsilon$ and that $W^* \geq \theta^{k_1}$. Since a new cut will be generated at u^{k_1}, we will have

$$W\left(u^{k_1}\right) + \left\langle g^{k_1}, u - u^{k_1}\right\rangle \leq W^{k_1}(u) \leq W^{k_2-1}(u), \quad \forall u \in \mathcal{U},$$

where g^{k_1} is a element of $\partial W(u^{k_1})$. In particular, choosing $u = u^{k_2}$ gives

$$W\left(u^{k_1}\right) + \left\langle g^{k_1}, u^{k_2} - u^{k_1}\right\rangle \leq W^{k_1}\left(u^{k_2}\right) \leq W^{k_2-1}\left(u^{k_2}\right) = \theta^{k_2-1} \leq W^*, \quad \forall u \in \mathcal{U}.$$

But $\varepsilon < W(u^{k_2}) - W^*$, so

$$\varepsilon < W\left(u^{k_2}\right) - W\left(u^{k_1}\right) - \left\langle g^{k_1}, u^{k_2} - u^{k_1}\right\rangle,$$

and the convexity of W gives

$$W\left(u^{k_2}\right) - W\left(u^{k_1}\right) \leq \left\langle g^{k_2}, u^{k_2} - u^{k_1}\right\rangle.$$

Therefore, since W has uniformly bounded subgradients, there exists $\kappa > 0$ such that

$$\varepsilon < 2\kappa \left\| u^{k_2} - u^{k_1} \right\|, \qquad \forall k_1, k_2 \in K_\varepsilon, k_1 \neq k_2.$$

Because $\mathcal{U}$ is compact, K_ε has to be finite. Otherwise there would exist a convergent subsequence of $\left\{u^k\right\}_{k \in K_\varepsilon}$ and this last inequality could not hold for sufficiently large indices within K_ε. This proves the lemma.

Note that Lemma 1 does not imply that the sequence of iterates $(u^k)_{k\in\mathbb{N}}$ converges[1]. For instance, if the minimum of W is attained on a "flat" part (if W is not strictly convex), then the sequence of iterates won't converge. However, the lemma shows that the sequence of W values at these iterates will converge.

We now consider the multi-stage case. Let T be a positive integer. We consider the following deterministic optimal control problem.

$$\begin{aligned}
\min_{x,u} \quad & \sum_{t=0}^{T-1} C_t(x_t, u_t) + V_T(x_T) && \text{(1a)}\\
\text{s.t.} \quad & x_{t+1} = f_t(x_t, u_t), \qquad \forall t = 0, \ldots, T-1, && \text{(1b)}\\
& x_0 \text{ is given}, && \text{(1c)}\\
& x_t \in \mathcal{X}_t, \qquad \forall t = 0, \ldots, T, && \text{(1d)}\\
& u_t \in \mathcal{U}_t, \qquad \forall t = 0, \ldots, T-1. && \text{(1e)}
\end{aligned}$$

Functions $C_t, t = 0, \ldots, T-1$ and V_T are assumed to be convex, as well as sets $\mathcal{X}_t \subseteq \mathbb{R}^n, t = 0, \ldots, T$ and $\mathcal{U}_t \subseteq \mathbb{R}^m, t = 0, \ldots, T-1$. Finally, the functions $f_t, t = 0, \ldots, T-1$ are linear so that (1) is a convex optimization problem.

The Dynamic Programming (DP) equation associated with this problem reads as follows. For all $t = 0, \ldots, T-1$:

$$V_t(x_t) = \min_{u_t \in \mathcal{U}_t} \underbrace{C_t(x_t, u_t) + V_{t+1}(f_t(x_t, u_t))}_{:= W_t(x_t, u_t)}, \quad \forall x_t \in \mathcal{X}_t. \tag{2}$$

[1] even though because $\mathcal{U}$ is compact, there exists a convergent subsequence

Quantity $W_t(x_t, u_t)$ is the future optimal cost starting at time t from state x and choosing decision u_t, so that $V_t(x) = \min_{u \in \mathcal{U}_t} W_t(x, u)$. We require an assumption that the future cost functions $V_t(x_t)$ at each stage have uniformly bounded subgradients, so

$$V_t(x) - V_t(y) \leq \gamma \|x - y\|, \quad x, y \in \mathbb{R}^n, \quad t = 0, \ldots, T. \tag{3}$$

This would appear to be a restrictive condition, but arises naturally if

$$f_t(x, \mathcal{U}_t) \cap \mathcal{X}_{t+1} \neq \emptyset$$

for every $x \in \mathcal{X}_t$. In other words, suppose we require that there exists a control at all times that maintains the feasibility of the states. Then this is ensured in practice by adding artificial variables v_t and w_t with high costs $(Mv_t + Mw_t)$ to each stage problem, to give

$$x_{t+1} = f_t(x_t, u_t) + v_t - w_t.$$

(We assume that there exists M so that $v_t = w_t = 0$ in an optimal solution.) This construction will prevent excursions of $f_t(x_t, u_t)$ outside $\mathcal{X}_{t+1}$, and provide a bound M on the Lagrange multipliers of this constraint, thereby bounding the subgradients of the future cost functions.

The cutting plane method works as follows. At iteration 0, define functions V_t^0, $t = 0, \ldots, T-1$, to be identically equal to $-\infty$. At time T, since we know exactly the end value function, we impose $V_T^k = V_T$ for all iterations $k \in \mathbb{N}$. At each iteration k, the process is the following. Starting with $x_0^k = x_0$, for $t = 0, 1, \ldots, T-1$, solve

$$\theta_t^k = \min_{u_t \in \mathcal{U}_t} \quad C_t(x, u_t) + V_{t+1}^{k-1} \circ f_t(x, u_t), \tag{4a}$$

$$\text{s.t.} \quad x = x_t^k. \qquad [\beta_t^k] \tag{4b}$$

Here β_t^k is the vector of Lagrange multipliers for the constraint $x = x_t^k$, and we denote the minimizer of this problem by u_t^k. Now define, for any $x \in \mathcal{X}_t$:

$$V_t^k(x) := \max\left(V_t^{k-1}(x), \theta_t^k + \left\langle \beta_t^k, x - x_t^k \right\rangle\right), \tag{5}$$

and $x_{t+1}^k = f_t(x_t^k, u_t^k)$.

Note that only the last future cost function V_T is known exactly at any iteration. All the other ones are lower approximations consisting of the supremum of k affine functions at iteration k. We naturally have the same lower approximation for function W_t, for any (x, u) in $\mathcal{X}_t \times \mathcal{U}_t$. Thus we define

$$W_t^k(x, u) := C_t(x, u) + V_{t+1}^k \circ f_t(x, u), \tag{6}$$

and recall

$$W_t(x, u) := C_t(x, u) + V_{t+1} \circ f_t(x, u). \tag{7}$$

Using this notation we have

$$\theta_t^k = \min_{u \in \mathcal{U}_t} W_t^{k-1}\left(x_t^k, u\right) = W_t^{k-1}\left(x_t^k, u_t^k\right) \tag{8}$$

Because for any $x \in \mathcal{X}_{t+1}$ and any $k' \leq k$ we have that $V_{t+1}^{k-1}(x) \geq V_{t+1}^{k'-1}(x)$, it follows that

$$\begin{aligned}
\theta_t^k = \min_{u \in \mathcal{U}_t} W_t^{k-1}\left(x_t^k, u\right) &\geq \min_{u \in \mathcal{U}_t} W_t^{k'-1}\left(x_t^k, u\right), \qquad \forall k' \leq k, \\
&= \min_{u \in \mathcal{U}_t} W_t^{k'-1}\left(x_t^{k'} + \left(x_t^k - x_t^{k'}\right), u\right), \qquad \forall k' \leq k,
\end{aligned}$$

which, using convexity of the optimal value function, and the definitions of $\theta_t^{k'}$ and $\beta_t^{k'}$, gives

$$\theta_t^k \geq \theta_t^{k'} + \left\langle \beta_t^{k'}, x_t^k - x_t^{k'} \right\rangle, \qquad \forall k' \leq k.$$

Since by (5)

$$V_t^{k-1}\left(x_t^k\right) = \max_{k'<k} \left\{ \theta_t^{k'} + \left\langle \beta_t^{k'}, x_t^k - x_t^{k'} \right\rangle \right\}$$

it follows that

$$\theta_t^k \geq V_t^{k-1}\left(x_t^k\right)$$

and so (5) implies

$$V_t^k\left(x_t^k\right) = \max\left(V_t^{k-1}\left(x_t^k\right), \theta_t^k\right) = \theta_t^k = W_t^{k-1}\left(x_t^k, u_t^k\right) \tag{9}$$

Figure 1 gives a view of the relations between all these values at a given iteration.

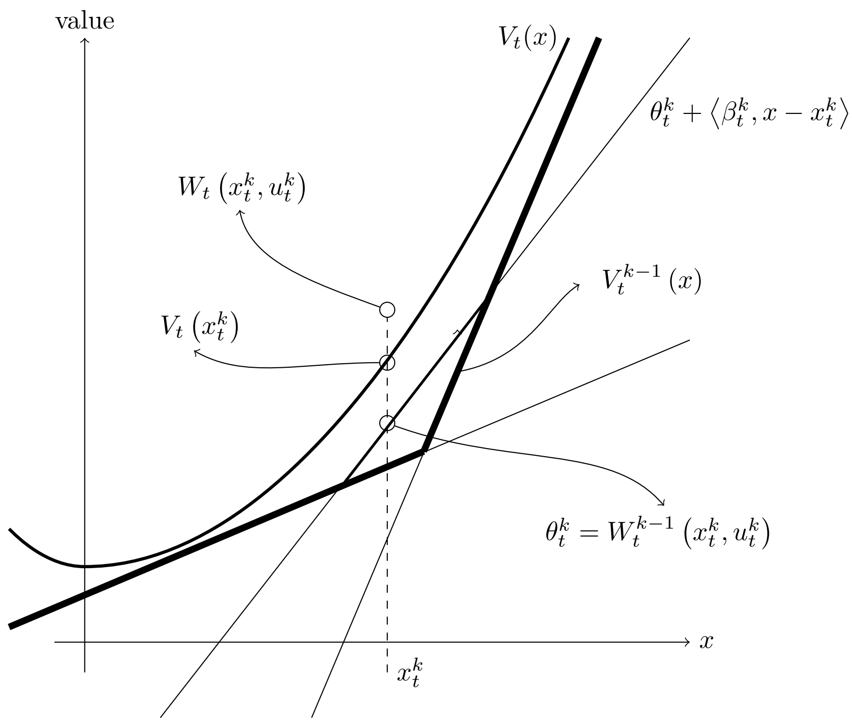

FIGURE 1. Relations between notations in the multi-stage case

We now prove that both the upper and lower estimates converge to the exact value function.

THEOREM 1. *Suppose that all future cost functions are convex and satisfy (3). Consider the sequence of decisions $(u^k)_{k\in\mathbb{N}}$ generated by the above described procedure, where each u^k is itself a sequence of decisions in time $u^k = u_0^k, \ldots, u_{T-1}^k$, and consider the corresponding sequence of state values $(x^k)_{k\in\mathbb{N}}$. Then for any time step $t = 0, \ldots, T-1$ we have that:*

$$\lim_{k\to+\infty} W_t\left(x_t^k, u_t^k\right) - V_t\left(x_t^k\right) = 0 \text{ and } \lim_{k\to+\infty} V_t\left(x_t^k\right) - V_t^k\left(x_t^k\right) = 0.$$

T he demonstration proceeds by induction backwards in time. At time $t+1$, the induction hypothesis is the second statement of the theorem. That is,

$$\lim_{k\to+\infty} V_{t+1}\left(x_{t+1}^k\right) - V_{t+1}^k\left(x_{t+1}^k\right) = 0.$$

In other words the cuts for the future cost function tend to be exact at x_{t+1}^k as k tends to ∞. The induction hypothesis is clearly true at the last time stage T for which we defined the approximate value function V_T^k to be equal to the (known) end value function V_T.

Now let us consider time t and choose K_ε to be the set of indices k such that:
1. $V_t\left(x_t^k\right) + \varepsilon < W_t\left(x_t^k, u_t^k\right) < +\infty$,
2. $V_{t+1}\left(x_{t+1}^k\right) - \theta_{t+1}^k < \frac{\varepsilon}{2}$.

Note that the induction hypothesis states

$$\lim_{k\to+\infty} V_{t+1}\left(x_{t+1}^k\right) - V_{t+1}^k\left(x_{t+1}^k\right) = 0$$

and (9) gives

$$V_{t+1}^k\left(x_{t+1}^k\right) = \theta_{t+1}^k$$

so

$$\lim_{k\to+\infty} V_{t+1}\left(x_{t+1}^k\right) - \theta_{t+1}^k = 0.$$

Hence the number of indices verifying condition (2) is infinite and if K_ε is finite, then it is the number of indices verifying condition (1) that has to be finite.

Let k_1, k_2 be two indices in K_ε, k_1 being smaller than k_2. Suppose that a new cut has just been added to the approximate future cost function at time $t+1$. Because of (5), we have that for any $k \in \mathbb{N}$ and any $x \in \mathcal{X}_{t+1}$

$$V_{t+1}^k(x) \ge \theta_{t+1}^k + \left\langle \beta_{t+1}^k, x - x_{t+1}^k \right\rangle.$$

It follows that the approximate future cost evaluated at $(x_t^{k_2}, u_t^{k_2})$ satisfies

$$\begin{aligned} C_t\left(x_t^{k_2}, u_t^{k_2}\right) + \theta_{t+1}^{k_1} + \left\langle \beta_{t+1}^{k_1}, x_{t+1}^{k_2} - x_{t+1}^{k_1} \right\rangle &\le C_t\left(x_t^{k_2}, u_t^{k_2}\right) + V_{t+1}^{k_1}\left(x_{t+1}^{k_2}\right), \\ &= W_t^{k_1}\left(x_t^{k_2}, u_t^{k_2}\right), \end{aligned}$$

from (6). Now, because $k_1 < k_2$,

$$W_t^{k_1}\left(x_t^{k_2}, u_t^{k_2}\right) \le W_t^{k_2-1}\left(x_t^{k_2}, u_t^{k_2}\right),$$

so (4) and (9) give

$$\begin{aligned} C_t\left(x_t^{k_2}, u_t^{k_2}\right) + \theta_{t+1}^{k_1} + \left\langle \beta_{t+1}^{k_1}, x_{t+1}^{k_2} - x_{t+1}^{k_1} \right\rangle &\le W_t^{k_2-1}\left(x_t^{k_2}, u_t^{k_2}\right) \\ &= \min_{u_t} W_t^{k_2-1}\left(x_t^{k_2}, u_t\right) \\ &\le \min_{u_t} W_t\left(x_t^{k_2}, u_t\right) \\ &= V_t\left(x_t^{k_2}\right). \end{aligned}$$

Now, using the fact that $k_2 \in K_\varepsilon$, we have

$$V_t\left(x_t^{k_2}\right) + \varepsilon < W_t\left(x_t^{k_2}, u_t^{k_2}\right)$$

which implies

$$C_t\left(x_t^{k_2}, u_t^{k_2}\right) + \theta_{t+1}^{k_1} + \left\langle \beta_{t+1}^{k_1}, x_{t+1}^{k_2} - x_{t+1}^{k_1} \right\rangle < W_t\left(x_t^{k_2}, u_t^{k_2}\right) - \varepsilon.$$

Now, the definition (7) implies

$$W_t\left(x_t^{k_2}, u_t^{k_2}\right) = C_t\left(x_t^{k_2}, u_t^{k_2}\right) + V_{t+1}\left(x_{t+1}^{k_2}\right)$$

so

$$\varepsilon < V_{t+1}\left(x_{t+1}^{k_2}\right) - \theta_{t+1}^{k_1} - \left\langle \beta_{t+1}^{k_1}, x_{t+1}^{k_2} - x_{t+1}^{k_1} \right\rangle,$$

or

$$\varepsilon < V_{t+1}\left(x_{t+1}^{k_1}\right) - \theta_{t+1}^{k_1} + V_{t+1}\left(x_{t+1}^{k_2}\right) - V_{t+1}\left(x_{t+1}^{k_1}\right) - \left\langle \beta_{t+1}^{k_1}, x_{t+1}^{k_2} - x_{t+1}^{k_1} \right\rangle \tag{10}$$

Now, $k_1 \in K_\varepsilon$ implies

$$V_{t+1}\left(x_{t+1}^{k_1}\right) - \theta_{t+1}^{k_1} < \frac{\varepsilon}{2},$$

and by (3) there is some γ_1 with

$$V_{t+1}\left(x_{t+1}^{k_2}\right) - V_{t+1}\left(x_{t+1}^{k_1}\right) \leq \gamma_1 \left\| x_{t+1}^{k_1} - x_{t+1}^{k_2} \right\|.$$

Futhermore, using the induction hypothesis and Lemma 2, we know that there exists a sequence of subgradients g_{t+1}^k of V_{t+1} at x_{t+1}^k such that:

$$\beta_{t+1}^k - g_{t+1}^k \xrightarrow{k\to\infty} 0.$$

By virtue of (3), the sequence $\left(\beta_{t+1}^k\right)_{k\in\mathbb{N}}$, made of the subgradients of lower approximations to V_{t+1}, is bounded as well. So there exists a constant $\gamma_2 > 0$ such that:

$$-\left\langle \beta_{t+1}^{k_1}, x_{t+1}^{k_2} - x_{t+1}^{k_1} \right\rangle \leq \gamma_2 \left\| x_{t+1}^{k_1} - x_{t+1}^{k_2} \right\|.$$

Finally, taking $\gamma = \gamma_1 + \gamma_2$ we obtain

$$\varepsilon < \frac{\varepsilon}{2} + \gamma \left\| x_{t+1}^{k_2} - x_{t+1}^{k_1} \right\|,$$

giving

$$\frac{\varepsilon}{2} < \gamma \left\| x_{t+1}^{k_2} - x_{t+1}^{k_1} \right\|. \tag{11}$$

It follows that K_ε has to be finite. Otherwise, since the state space $\mathcal{X}_{t+1}$ is compact, we could find a convergent subsequence of states verifying (11), and passing to the limit would lead to $\frac{\varepsilon}{2} \leq 0$.

We now have to show the induction hypothesis, namely

$$\lim_{k\to+\infty} V_t\left(x_t^k\right) - V_t^k\left(x_t^k\right) = 0$$

for time t. Recall (9) gives

$$\begin{aligned} V_t^k\left(x_t^k\right) = \theta_t^k = W_t^{k-1}\left(x_t^k, u_t^k\right), \\ = C_t\left(x_t^k, u_t^k\right) + V_{t+1}^{k-1}\left(x_{t+1}^k\right). \end{aligned}$$

Using the definition (7) of W_t, we can replace $C_t\left(x_t^k, u_t^k\right)$ to get

$$V_t^k\left(x_t^k\right) = W_t\left(x_t^k, u_t^k\right) + \left(V_{t+1}^{k-1}\left(x_{t+1}^k\right) - V_{t+1}\left(x_{t+1}^k\right)\right).$$

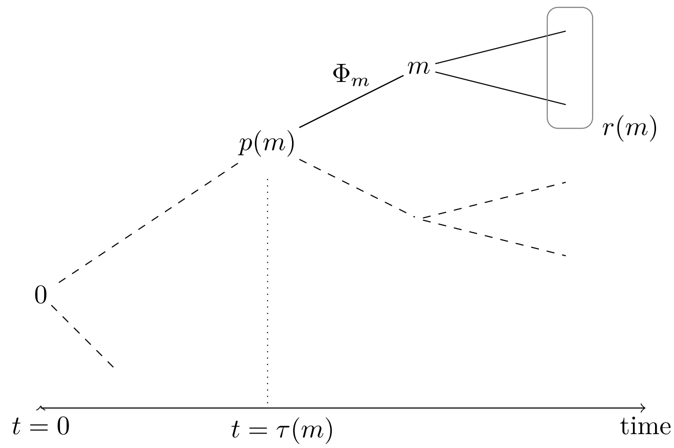

FIGURE 2. Notations on the scenario tree

We have just demonstrated above that

$$\delta_t^k = W_t\left(x_t^k, u_t^k\right) - V_t\left(x_t^k\right) \xrightarrow{k\to\infty} 0$$

so

$$V_t^k\left(x_t^k\right) = V_t\left(x_t^k\right) + \delta_t^k + \left(V_{t+1}^{k-1}\left(x_{t+1}^k\right) - V_{t+1}\left(x_{t+1}^k\right)\right), \text{ with } \delta_t^k \xrightarrow{k\to\infty} 0.$$

The induction hypothesis at time $t+1$ gives

$$V_{t+1}^k\left(x_{t+1}^k\right) - V_{t+1}\left(x_{t+1}^k\right) \xrightarrow{k\to\infty} 0,$$

which by virtue of Lemma 3 (with V_{t+1} replacing f) implies

$$\lim_{k\to+\infty} V_{t+1}^k\left(x_{t+1}^k\right) - \lim_{k\to+\infty} V_{t+1}^{k-1}\left(x_{t+1}^k\right) = 0$$

so

$$V_t^k\left(x_t^k\right) - V_t\left(x_t^k\right) \xrightarrow{k\to\infty} 0$$

which gives the result.

Theorem 1 indicates that the lower approximation at each iteration tends to be exact on the sequence of state trajectories generated throughout the algorithm. This does not mean that the future cost function will be approximated well everywhere in the state space. It only means that the approximation gets better and better in the neighbourhood of an optimal state trajectory.

3. The stochastic case with a finite distribution We now consider that the cost function and dynamics at each time t are influenced by a random outcome that has a discrete and finite distribution. We write the problem on the complete tree induced by this distribution. The set of all nodes is denoted by $\mathcal{N}$ and $\{0\}$ is the root node. We denote nodes by m and n. (We trust that the context will dispel any confusion from the use of m and n as dimensions of variables u and x.) For every node $m \in \mathcal{N}\backslash\{0\}$, $p(m)$ represents its parent, $r(m)$ its set of children nodes, and Φ_m its probability of occurence. Finally, $\tau(m)$ is the time step at which node m appears, and $\mathcal{L} = \{m \mid \tau(m) = T\}$ is the set of leaves. Figure 2 illustrates this notation.

This gives the following stochastic program:

$$\min_{x,u} \quad \sum_{m\in\mathcal{N}\backslash\{0\}} \Phi_m \cdot C_m\left(x_{p(m)}, u_m\right) + \sum_{m\in\mathcal{L}} \Phi_n V_m\left(x_m\right) \tag{12a}$$

$$\text{s.t.} \quad x_m = f_m\left(x_{p(m)}, u_m\right), \qquad \forall m \in \mathcal{N}\backslash\{0\}, \tag{12b}$$

$$x_0 \text{ is given}, \tag{12c}$$

$$x_m \in \mathcal{X}_{\tau(m)}, \qquad \forall m \in \mathcal{N}, \tag{12d}$$

$$u_m \in \mathcal{U}_{\tau(m)-1}, \qquad \forall m \in \mathcal{N}\backslash\{0\}. \tag{12e}$$

The reader should note that randomness (that appears in the cost and in the dynamics) occurs before the decision is taken in this model. Hence the control affecting the stock[2] x_n is actually indexed by m, a child node of n. Put differently, the control adapts to randomness: there are as many controls as there are children nodes of n.

We have a future cost function for each node $n \in \mathcal{N}$ which is defined by

$$V_n(x_n) = \sum_{m \in r(n)} \frac{\Phi_m}{\Phi_n} \min_{u_m \in \mathcal{U}_t} \underbrace{C_m(x_n, u_m) + V_m(f_m(x_n, u_m))}_{W_m(x_n, u_m)}. \tag{13}$$

As in the previous section we require that there is some $\gamma > 0$ with

$$V_m(x) - V_m(y) \leq \gamma \|x - y\|, \quad x, y \in \mathbb{R}^n, \quad m \in \mathcal{N}. \tag{14}$$

In general the future cost function at each node is different from those at other nodes at the same stage. The algorithm that we consider is now an extension of the deterministic algorithm of the previous section applied, at each iteration, to a set of nodes chosen randomly in the tree to update estimates of the future cost function at these nodes. To help the exposition we assume that all other nodes have null updates, in the sense that they just inherit the future cost function from the previous iteration.

We now describe the algorithm formally. We start the process with $\hat{\theta}_n^0 = -\infty$, $\hat{\beta}_n^0 = 0$, for each $n \in \mathcal{N}$, and impose[3] $V_n^k = V_n$ for all nodes $n \in \mathcal{L}$ and all $k \in \mathbb{N}$. We then carry out a sequence of simulations and updates of the future cost functions as follows.
Simulation Starting at the root node, generate states and decisions for all possible successors (in other words, visit the whole tree forward) by solving (13) with V^{k-1} instead of V. Denote the obtained state variables by $(x_n^k)_{n \in \mathcal{N}}$ and the control variables by $(u_n^k)_{n \in \mathcal{N} \setminus \{0\}}$. Also, for each node $n \in \mathcal{N}$, impose $\theta_n^k = V_n^{k-1}(x_n^k)$ and $\beta_n^k \in \partial V_n^{k-1}(x_n^k)$.
Update Select non-terminal nodes $n_1, n_2, \ldots, n_I$ in the tree. For each $i = 1, 2, \ldots, I$, $x_{n_i}^k$ is hence a random variable which is equal to one of the x_n^k. For each selected node n_i, and for every child node m of node n_i, solve:

$$\begin{aligned} \hat{\theta}_m^k = \min_{u_m \in \mathcal{U}_{T-1}} \; & C_m(x, u_m) + V_m^{k-1} \circ f_m(x, u_m), \\ \text{s.t.} \quad & x = x_{n_i}^k. \qquad [\hat{\beta}_m^k] \end{aligned}$$

As before $\hat{\beta}_m^k$ is the Lagrange multiplier at optimality. For each selected node n_i, replace the values $\theta_{n_i}^k$ and $\beta_{n_i}^k$ obtained during the forward pass with $\theta_{n_i}^k = \sum_{m \in r(n_i)} \frac{\Phi_m}{\Phi_{n_i}} \hat{\theta}_m^k$ and $\beta_{n_i}^k = \sum_{m \in r(n_i)} \frac{\Phi_m}{\Phi_{n_i}} \hat{\beta}_m$.

Finally, we update all future cost functions. For every node n

$$V_n^k(x) := \max\left(V_n^{k-1}(x), \theta_n^k + \left\langle \beta_n^k, x - x_n^k \right\rangle\right), \quad x \in \mathcal{X}_t. \tag{15}$$

Note that we actually only update future cost functions on the selected nodes. Since the cuts we add at all other nodes are binding on the current model (by construction in the forward pass), there is no point in storing them. Therefore, in practice, one does not need to sample the whole scenario tree but just enough to attain all selected nodes. In our proof, we need to look at what happens even on the nodes that are not selected.

[2] We do not make any stagewise independence assumptions on the random variables that affect the system. Hence there is no reason why variable x_n should be called a state variable and we prefer calling it a stock.

[3] Because the end value function is known exactly, there is no need to approximate it.

The way $n_1, n_2, \ldots, n_I$ are selected is actually irrelevant, provided that all nodes in the tree are visited randomly with positive probabilities in a way that, at every iteration, is independent of the current values of the decision variables on the tree. Let us explain this more formally. Denote by y_n^k the random variable that is equal to 1 if node n is selected at iteration k and 0 otherwise. We require that for any $n \in \mathcal{N}$ and any $k \in \mathbb{N}$, the random variable y_n^k is independent from the current error $V_n\left(x_n^k\right) - V_n^k\left(x_n^k\right)$. Otherwise, it would be easy to create a case in which the future cost function at a given node is updated only when the current stock variable on this node is in a given region for instance. In such a case the future cost function could not gather any information about the other parts of the space that the stock variable might visit. In other words, this independence assumption ensures that the values that are optimal can be attained an infinite number of times.

We will make use of the following definitions:

$$W_m\left(x_n, u_m\right) := C_m\left(x_n, u_m\right) + V_m\left(f_m\left(x_n, u_m\right)\right) \tag{16}$$

$$W_m^k\left(x_n, u_m\right) := C_m\left(x_n, u_m\right) + V_m^k\left(f_m\left(x_n, u_m\right)\right) \tag{17}$$

In the case where node $n \in \mathcal{N}$ is selected at iteration k, in other words $n = n_i$, these definitions then give

$$\hat{\theta}_m^k = \min_{u \in \mathcal{U}_t} W_m^{k-1}\left(x_n^k, u\right) = W_m^{k-1}\left(x_n^k, u_m^k\right).$$

Because for any $x \in \mathcal{X}_{t+1}$ and any $k' \leq k$ we have that $V_m^{k'-1}(x) \leq V_m^{k-1}(x)$, it follows that

$$\begin{aligned}\hat{\theta}_m^k = \min_{u \in \mathcal{U}_t} W_m^{k-1}\left(x_n^k, u\right) &\geq \min_{u \in \mathcal{U}_t} W_m^{k'-1}\left(x_n^k, u\right), \qquad \forall k' \leq k, \\ &= \min_{u \in \mathcal{U}_t} W_m^{k'-1}\left(x_n^{k'} + \left(x_n^k - x_n^{k'}\right), u\right), \qquad \forall k' \leq k,\end{aligned}$$

which, using convexity of the value function of this latter problem and the definitions of $\hat{\theta}_m^{k'}$ and $\hat{\beta}_m^{k'}$, gives

$$\hat{\theta}_m^k \geq \hat{\theta}_m^{k'} + \left\langle \hat{\beta}_m^{k'}, x_n^k - x_n^{k'} \right\rangle, \qquad \forall k' \leq k.$$

and taking conditional expectations with probabilities $\frac{\Phi_m}{\Phi_n}$, we obtain

$$\theta_n^k \geq \theta_n^{k'} + \left\langle \beta_n^{k'}, x_n^k - x_n^{k'} \right\rangle, \qquad \forall k' \leq k.$$

Since by (15)

$$V_n^{k-1}\left(x_n^k\right) = \max_{k' < k} \{\theta_n^{k'} + \left\langle \beta_n^{k'}, x_n^k - x_n^{k'} \right\rangle\}$$

it follows that

$$\theta_n^k \geq V_n^{k-1}\left(x_t^k\right),$$

and so (15) implies (once again, only in the case when $n = n_i$)

$$V_n^k\left(x_n^k\right) = \max\left(V_n^{k-1}\left(x_n^k\right), \theta_n^k\right) = \theta_n^k = \sum_{m \in r(n)} \frac{\Phi_m}{\Phi_n} W_m^{k-1}\left(x_n^k, u_m^k\right). \tag{18}$$

Note that because of the way x_n^k and θ_n^k are defined during simulation for the non-selected nodes, we have

$$V_n^k\left(x_n^k\right) = \max\left(V_n^{k-1}\left(x_n^k\right), \theta_n^k\right) = \theta_n^k, \text{ for all } n \in \mathcal{N}.$$

THEOREM 2. *Suppose that all future cost functions are convex with bounded subgradients. Consider the sequence of decisions $(u^k)_{k\in\mathbb{N}}$ generated by the above described procedure, where each u^k is itself a set of decisions on the complete tree $u^k = u_1^k, \dots, u_{N-1}^k$, and consider the corresponding sequence of state values $(x^k)_{k\in\mathbb{N}}$. Assume random variables $y_n^k, k \in \mathbb{N}$, are i.i.d. with $\mathbb{P}(y_n^k = 1) = p_n > 0$ and that for any $n \in \mathcal{N}$ and any $k \in \mathbb{N}$ the selection variable y_n^k is independent from $V_n(x_n^k) - V_n^{k-1}(x_n^k)$. Then for any node $n = 0, \dots, N-1$ we have that, $\mathbb{P}$-almost surely:*

$$\lim_{k\to+\infty} W_n\left(x_n^k, u_n^k\right) - V_n\left(x_n^k\right) = 0 \text{ and } \lim_{k\to+\infty} V_n\left(x_n^k\right) - V_n^k\left(x_n^k\right) = 0.$$

B ecause the selection process for nodes in the update step is stochastic, decision variables as well as approximate future cost functions are stochastic throughout the course of the algorithm. Thus, during the whole proof, all equalities or inequalities are taken $\mathbb{P}$-almost surely.

The demonstration follows the same approach as the proof of Theorem 1. At time $t+1$, the induction hypothesis is

$$\lim_{k\to+\infty} V_m\left(x_m^k\right) - V_m^k\left(x_m^k\right) = 0$$

for each node m with $\tau(m) = t+1$. Since at the very last time stage those two quantities are equal, by definition, the induction hypothesis is true for every node n with $\tau(n) = T$. Let us consider an arbitrary node n with $\tau(m) = t$, and choose K_ε to be the set of indices k such that:

1. $V_n\left(x_n^k\right) + \varepsilon < \sum_{m\in r(n)} \frac{\Phi_m}{\Phi_n} W_m\left(x_n^k, u_m^k\right) < +\infty$,
2. $\sum_{m\in r(n)} \frac{\Phi_m}{\Phi_n}\left(V_m\left(x_m^k\right) - \theta_m^k\right) < \frac{\varepsilon}{2}$.

Note that because of the induction hypothesis and (18) the quantity on the left-hand side of condition (2) converges to 0 as k goes to infinity. Hence the number of indices verifying condition (2) is infinite and if K_ε is finite, then it is the number of indices verifying condition (1) that has to be finite.

Because of (15), we have that for any $k \in \mathbb{N}$, any $x \in \mathcal{X}_{t+1}$ and any child node m of n:

$$V_m^k(x) \geq \theta_m^k + \left\langle \beta_m^k, x - x_m^k \right\rangle. \tag{19}$$

Let k_1, k_2 be two indices in K_ε, k_1 being smaller than k_2. It follows from (19) at $k = k_1$, that the approximate future cost evaluated at $(x_n^{k_2}, u_m^{k_2})$ satisfies

$$\begin{aligned} C_m\left(x_n^{k_2}, u_m^{k_2}\right) + \theta_m^{k_1} + \left\langle \beta_m^{k_1}, x_m^{k_2} - x_m^{k_1} \right\rangle &\leq C_m\left(x_n^{k_2}, u_m^{k_2}\right) + V_m^{k_1}\left(x_m^{k_2}\right), \\ &= W_m^{k_1}\left(x_n^{k_2}, u_m^{k_2}\right). \end{aligned}$$

and $k_1 < k_2$ implies

$$W_m^{k_1}\left(x_n^{k_2}, u_m^{k_2}\right) \leq W_m^{k_2-1}\left(x_n^{k_2}, u_m^{k_2}\right).$$

Thus

$$\begin{aligned} \sum_{m\in r(n)} \frac{\Phi_m}{\Phi_n}\Big(C_m\left(x_n^{k_2}, u_n^{k_2}\right) + \theta_m^{k_1} + \left\langle \beta_m^{k_1}, x_m^{k_2} - x_m^{k_1} \right\rangle\Big) &\leq \sum_{m\in r(n)} \frac{\Phi_m}{\Phi_n} W_m^{k_2-1}\left(x_n^{k_2}, u_m^{k_2}\right), \\ &= \sum_{m\in r(n)} \frac{\Phi_m}{\Phi_n} \min_{u_m} W_m^{k_2-1}\left(x_n^{k_2}, u_m\right), \end{aligned}$$

by definition of $u_m^{k_2}$ (simulation step)

$$\leq \sum_{m\in r(n)} \frac{\Phi_m}{\Phi_n} \min_{u_m} W_m\left(x_n^{k_2}, u_m\right),$$

because $V_m^{k_2-1}$ is identically lower than V_m

$$= V_n\left(x_n^{k_2}\right).$$

Now, since $k_2 \in K_\varepsilon$,

$$V_n\left(x_n^{k_2}\right) + \varepsilon < \sum_{m\in r(n)} \frac{\Phi_m}{\Phi_n} W_m\left(x_n^{k_2}, u_m^{k_2}\right)$$

which gives

$$\sum_{m\in r(n)} \frac{\Phi_m}{\Phi_n}\left(C_m\left(x_n^{k_2}, u_n^{k_2}\right) + \theta_m^{k_1} + \left\langle \beta_m^{k_1}, x_m^{k_2} - x_m^{k_1}\right\rangle\right) < \sum_{m\in r(n)} \frac{\Phi_m}{\Phi_n} W_m\left(x_n^{k_2}, u_m^{k_2}\right) - \varepsilon.$$

Now recall (16) the definition of W_m which gives

$$\sum_{m\in r(n)} \frac{\Phi_m}{\Phi_n} W_m\left(x_n^{k_2}, u_m^{k_2}\right) = \sum_{m\in r(n)} \frac{\Phi_m}{\Phi_n} C_m\left(x_n^{k_2}, u_m^{k_2}\right) + \sum_{m\in r(n)} \frac{\Phi_m}{\Phi_n} V_m\left(f_m\left(x_n^{k_2}, u_m^{k_2}\right)\right).$$

Substituting and observing

$$x_m^{k_2} = f_m\left(x_n^{k_2}, u_m^{k_2}\right)$$

gives

$$\varepsilon < \sum_{m\in r(n)} \frac{\Phi_m}{\Phi_n}\left(V_m\left(x_m^{k_2}\right) - \theta_m^{k_1} - \left\langle \beta_m^{k_1}, x_m^{k_2} - x_m^{k_1}\right\rangle\right),$$

yielding

$$\varepsilon < \sum_{m\in r(n)} \frac{\Phi_m}{\Phi_n}\left(V_m\left(x_m^{k_1}\right) - \theta_m^{k_1}\right) + \sum_{m\in r(n)} \frac{\Phi_m}{\Phi_n}\left(V_m\left(x_m^{k_2}\right) - V_m\left(x_m^{k_1}\right)\right) - \sum_{m\in r(n)} \frac{\Phi_m}{\Phi_n}\left\langle \beta_m^{k_1}, x_m^{k_2} - x_m^{k_1}\right\rangle. \tag{20}$$

Now, $k_1 \in K_\varepsilon$ implies

$$\sum_{m\in r(n)} \frac{\Phi_m}{\Phi_n}\left(V_m\left(x_m^{k_1}\right) - \theta_m^{k_1}\right) < \frac{\varepsilon}{2}.$$

Futhermore by (14) there exists a constant $\gamma_1 > 0$ such that

$$V_m\left(x_m^{k_2}\right) - V_m\left(x_m^{k_1}\right) \le \gamma_1 \left\|x_m^{k_1} - x_m^{k_2}\right\|.$$

Finally, using the induction hypothesis and Lemma 2, we know that there exists a sequence of subgradients g_m^k of V_m at x_m^k such that

$$\beta_m^k - g_m^k \xrightarrow{k\to\infty} 0.$$

Therefore, the sequence $(\beta_m^k)_{k\in\mathbb{N}}$, made of the subgradients of lower approximations to V_m, is bounded as well. So there exists a constant $\gamma_2 > 0$ such that:

$$-\left\langle \beta_m^{k_1}, x_m^{k_2} - x_m^{k_1}\right\rangle \le \gamma_2 \left\|x_m^{k_1} - x_m^{k_2}\right\|.$$

Collecting terms and taking $\gamma = \gamma_1 + \gamma_2$ we obtain

$$\varepsilon < \frac{\varepsilon}{2} + \gamma \sum_{m\in r(n)} \frac{\Phi_m}{\Phi_n}\left\|x_m^{k_2} - x_m^{k_1}\right\|,$$

giving

$$\frac{\varepsilon}{2} < \gamma \sum_{m \in r(n)} \frac{\Phi_m}{\Phi_n} \left\| x_m^{k_2} - x_m^{k_1} \right\|. \tag{21}$$

Hence K_ε has to be finite. Otherwise, because the space $(\mathcal{X}_{t+1})^N$ is compact, we could find a convergent subsequence of states verifying (21), and passing to the limit would lead to $\frac{\varepsilon}{2} \leq 0$.

We now have to show the induction hypothesis, namely

$$\lim_{k \to +\infty} V_n\left(x_n^k\right) - V_n^k\left(x_n^k\right) = 0$$

for every node n at time t. We start by proving it for iterations k such that n is selected, i.e. such that $y_n^k = 1$. Recall (18) for any selected node n at iteration k

$$V_n^k\left(x_n^k\right) = \max\left(V_n^{k-1}\left(x_n^k\right), \theta_n^k\right) = \theta_n^k = \sum_{m \in r(n)} \frac{\Phi_m}{\Phi_n} W_m^{k-1}\left(x_n^k, u_m^k\right)$$

which implies

$$\begin{aligned} V_n^k\left(x_n^k\right) &= \sum_{m \in r(n)} \frac{\Phi_m}{\Phi_n} \left[C_m\left(x_n^k, u_m^k\right) + V_m^{k-1}\left(x_m^k\right)\right], \\ &= \sum_{m \in r(n)} \frac{\Phi_m}{\Phi_n} \left[W_m\left(x_n^k, u_m^k\right) + \left(V_m^{k-1}\left(x_m^k\right) - V_m\left(x_m^k\right)\right)\right]. \end{aligned}$$

We have just demonstrated above that

$$\delta_n^k = \sum_{m \in r(n)} \frac{\Phi_m}{\Phi_n} W_m\left(x_n^k, u_m^k\right) - V_n\left(x_n^k\right) \xrightarrow{k \to \infty} 0.$$

Thus

$$V_n^k\left(x_n^k\right) = V_n\left(x_n^k\right) + \delta_n^k + \sum_{m \in r(n)} \frac{\Phi_m}{\Phi_n} \left(V_m^{k-1}\left(x_m^k\right) - V_m\left(x_m^k\right)\right)$$

with $\delta_n^k \xrightarrow{k \to \infty} 0$.

The inductive hypothesis gives

$$\lim_{k \to +\infty} V_m\left(x_m^k\right) - V_m^k\left(x_m^k\right) = 0$$

which by virtue of Lemma 3 (with V_m replacing f) implies

$$\lim_{k \to +\infty} V_m\left(x_m^k\right) - V_m^{k-1}\left(x_m^k\right) = 0.$$

Finally

$$V_n\left(x_n^k\right) - V_n^k\left(x_n^k\right) \xrightarrow[y_n^k = 1]{k \to \infty} 0. \tag{22}$$

Now we prove that the values also converge for the iterations k such that $y_n^k = 0$, i.e. the iterations for which node n is not selected. By contradiction, suppose the values do not converge. Then

$$\exists \varepsilon > 0, \forall K \in \mathbb{N}, \exists k \geq K, \quad V_n\left(x_n^k\right) - V_n^k\left(x_n^k\right) \geq \varepsilon.$$

and, because $V_n^{k-1}\left(x_n^k\right) \leq V_n^k\left(x_n^k\right)$, we have a fortiori that

$$\exists \varepsilon > 0, \forall K \in \mathbb{N}, \exists k \geq K, \quad V_n\left(x_n^k\right) - V_n^{k-1}\left(x_n^k\right) \geq \varepsilon. \tag{23}$$

We denote by $\mathcal{K}_\varepsilon$ the set of indices satisfying this relation. We note that:

- the set $\mathcal{K}_\varepsilon \cap \{k \in \mathbb{N} | y_n^k = 1\}$ is finite because of (22) and Lemma 3 (with V_n replacing f);
- the set $\mathcal{K}_\varepsilon \cap \{k \in \mathbb{N} | y_n^k = 0\}$ is infinite because of (23).

Let z^j denote the j-th element of the set $\{y_n^k, k \in \mathcal{K}_\varepsilon\}$. Note that for any k the quantity $V_n(x_n^k) - V_n^{k-1}(x_n^k)$ is independent from y_n^k and thus so is the event $\{k \in \mathcal{K}_\varepsilon\}$. This implies that random variables z^j are i.i.d. and share the same probability law as any of the y_n^k, $k \in \mathbb{N}$. According to the Strong Law of Large Numbers [5] applied to the random sequence (z^j), we should then have

$$\frac{1}{N}\sum_{j=1}^{N} z^j \xrightarrow{N\to+\infty} \mathbb{E}\left[y_n^k \mid k \in \mathcal{K}_\varepsilon\right] = \mathbb{E}\left[y_n^0\right] = p_n > 0.$$

However, because $\mathcal{K}_\varepsilon \cap \{y_n^k = 1\}$ is finite, we know that there is only a finite number of indices j such that $z^j = 1$, the rest being equal to 0. So

$$\frac{1}{N}\sum_{j=1}^{N} z^j \xrightarrow{N\to+\infty} 0,$$

which is a contradiction. This shows that

$$V_n\left(x_n^k\right) - V_n^k\left(x_n^k\right) \xrightarrow[y_n^k=0]{k\to\infty} 0.$$

and concludes the proof.

Remark 1. The theorem uses random sampling but one can observe that a procedure that uses a deterministic selection process satisfying

$$\frac{1}{N}\sum_{k=1}^{N} y_n^k \xrightarrow{N\to+\infty} p_n \tag{24}$$

will converge as well. (This will be the case for a round robin procedure, for example.) The only change needed in the proof of Theorem 2 is in the last paragraph where, instead of using the Strong Law of Large Numbers, we simply use Equation (24) and observe the same contradiction.

4. Discussion

The convergence result we proved is very general. We require an independence assumption in the selection of nodes at which to construct new cutting planes. The state values at which these are constructed are generated by simulation, but this simulation need not encompass every outcome.

Our convergence result applies to some well-known sampling-based algorithms for multistage stochastic programs. The CUPPS algorithm of [1] is easily seen to fall into our framework. Here cuts are computed at every node on a randomly chosen scenario at the states generated by incumbent policies.

The DOASA algorithm [9] adds cutting planes on a randomly chosen scenario, starting at the end of the time horizon and working backwards in time. So the approximations of the future cost functions are changed as we traverse this scenario in the backward pass. To place this algorithm in our setting, observe that a single backward pass can be seen as T simulation/update iterations of the procedure described on page 9. In other words, we start by simulating up to time stage $T-1$, compute a cutting plane for the selected node at this stage and update the approximate future cost function. Then we do the same for the node of the previous time stage, but observe it can be seen as a new simulation/update iteration where, because the future cost functions have not changed up to $T-1$, the simulated state values stay the same.

Now, SDDP, that uses several forward passes, can be viewed in a similar way. It simulates a fixed number N of scenarios, adds cutting planes for the N nodes selected at time $T-1$, updates the future cost functions there, and moves on to the previous stage. This can be seen as T iterations of the procedure described on page 9, where $n_1, n_2, \ldots, n_I$ are the N selected nodes of a given time stage.

For both of these algorithms (SDDP and DOASA), note that the independence assumption that concerns the selection process is fulfilled. As soon as the leaf node is selected, the previous nodes on the scenario, for the next $T-2$ iterations, are determined. However, the fact that this leaf node was selected independently of previous computations ensures independence between the selection of any node and previous computations at that node.

We have proved this result with a general scenario tree. In SDDP algorithms, the random variables are usually assumed to be stagewise independent (or made so by adding state variables). This means that the future cost functions $V_m(x)$ at each node m with $\tau(m) = t$ are the same. This allows cutting planes in the approximations to be shared across these nodes. The convergence result we have shown here applies to this situation as a special case.

Also note that, in the case where one adds cuts at different nodes in the tree in the update step of our procedure, the solving of the subproblems can be done in parallel. This is the case in CUPPS, where a whole branch of the tree is selected at each iteration. It also allows us to consider different selection strategies, where nodes at a given iteration could be selected throughout the tree depending on some criteria defined by the user. In the first few iterations, this could increase efficiency of the approximation and, because the solving of the subproblems can be parallelized, would not be very time-consuming. However, one should bear in mind that, at some point, the algorithm has to come back to an appropriate selection procedure, i.e. one that satisfies the independence assumption, in order to ensure convergence of the algorithm.

Appendix. Technical lemmas

LEMMA 2. *Let f and $f^k, k \in \mathbb{N}$, be convex functions on $\mathbb{R}^n$, such that functions f^k are identically lower than f. Suppose that there exists a sequence $(x^k)_{k\in\mathbb{N}}$ such that:*

$$f^k\left(x^k\right) - f\left(x^k\right) \xrightarrow{k\to\infty} 0.$$

Then, for any sequence $(\beta^k)_{k\in\mathbb{N}}$ of elements of $\partial f^k(x^k)$, there exists a sequence $(g^k)_{k\in\mathbb{N}}$ of elements of $\partial f(x^k)$ such that:

$$\beta^k - g^k \xrightarrow{k\to\infty} 0.$$

S uppose the result is not true. Then it means that we cannot construct such a sequence $(g^k)_{k\in\mathbb{N}}$. In other words, there is some $\varepsilon > 0$ and a subsequence $\left(\beta^{k(l)}\right)_{l\in\mathbb{N}}$ with $d\left(\beta^{k(l)}, \partial f\left(x^{k(l)}\right)\right) > \varepsilon$, where d is the Euclidean distance between a point and a set. Because a subdifferential is a convex set, it implies that $\beta^{k(l)}$ can be strictly separated from $\partial f\left(x^{k(l)}\right)$ using a hyperplane. There exists $c^l \in \mathbb{R}$ and a vector $\gamma^l \in \mathbb{R}^n$ such that

$$\left\langle \gamma^l, \beta^{k(l)} \right\rangle > c^l \text{ and } \left\langle \gamma^l, g \right\rangle + \varepsilon < c^l, \quad \forall g \in \partial f\left(x^{k(l)}\right).$$

Because the directional derivative $\nabla_{\gamma^l} f\left(x^{k(l)}\right) = \sup_{g\in\partial f\left(x^{k(l)}\right)} \left\langle \gamma^l, g \right\rangle$, we have that

$$\nabla_{\gamma^l} f\left(x^{k(l)}\right) + \varepsilon \le c^l < \left\langle \gamma^l, \beta^{k(l)} \right\rangle \le \nabla_{\gamma^l} f^{k(l)}\left(x^{k(l)}\right),$$

which, by definition of the directional derivative, indicates that

$$\exists \delta^l > 0, \forall t < \delta^l, f\left(x^{k(l)} + t\gamma^l\right) - f\left(x^{k(l)}\right) \le f^{k(l)}\left(x^{k(l)} + t\gamma^l\right) - f^{k(l)}\left(x^{k(l)}\right) - \varepsilon.$$

Now

$$f^k\left(x^k\right) - f\left(x^k\right) \xrightarrow{k\to\infty} 0.$$

so there exists $l > L$ with

$$f\left(x^{k(l)}\right) - f^{k(l)}\left(x^{k(l)}\right) < \frac{\varepsilon}{2},$$

and so:

$$f\left(x^{k(l)} + t\gamma^l\right) < f^{k(l)}\left(x^{k(l)} + t\gamma^l\right) - \frac{\varepsilon}{2},$$

which contradicts the fact that $f^{k(l)}$ is identically lower than f.

We will also need the following result.

Lemma 3. *Suppose f is convex with bounded subgradients and $\mathcal{X}$ is compact, and suppose the sequence of convex functions $f^k, k \in \mathbb{N}$, satisfies*

$$f^{k-1}(x) \le f^k(x) \le f(x)\,, \textit{ for all } x \in \mathcal{X}.$$

Then, for any infinite sequence $x^k \in \mathcal{X}$

$$\lim_{k\to+\infty} f\left(x^k\right) - f^k\left(x^k\right) = 0 \Rightarrow \lim_{k\to+\infty} f\left(x^k\right) - f^{k-1}\left(x^k\right) = 0.$$

A ssume

$$\lim_{k\to+\infty} f\left(x^k\right) - f^k\left(x^k\right) = 0$$

and suppose the result is wrong. Then

$$\exists\varepsilon > 0, \exists K \in \mathbb{N}, \text{such that } \forall k > K,\ f\left(x^k\right) - f^{k-1}\left(x^k\right) > \varepsilon.$$

Now, because $\mathcal{X}$ is compact, there exists a subsequence $\left(x^{k(l)}\right)_{l\in\mathbb{N}}$ that converges to $x^* \in \mathcal{X}$. For sufficiently large l, the continuity of $f^{k(l)}$ and $f^{k(l)-1}$ gives

$$\left|f^{k(l)}(x^*) - f^{k(l)}\left(x^{k(l)}\right)\right| < \frac{\varepsilon}{3},$$
$$\left|f^{k(l)-1}\left(x^{k(l)}\right) - f^{k(l)-1}(x^*)\right| < \frac{\varepsilon}{3},$$

so

$$\begin{aligned}
& f^{k(l)}(x^*) - f^{k(l)-1}(x^*) \\
= & f^{k(l)}(x^*) - f^{k(l)}\left(x^{k(l)}\right) \\
& + f^{k(l)}\left(x^{k(l)}\right) - f^{k(l)-1}\left(x^{k(l)}\right) \\
& + f^{k(l)-1}\left(x^{k(l)}\right) - f^{k(l)-1}(x^*) \\
> & \frac{\varepsilon}{3}
\end{aligned}$$

So $f^{k(l)}(x^*) > f^{k(l)-1}(x^*) + \frac{\varepsilon}{3}$, for infinitely many l which contradicts the fact that $f^k(x^*)$ is bounded above by $f(x^*)$.

Acknowledgments. The first author was supported in this research by a grant from the OSIRIS Department at EDF R&D, France. The authors also wish to thank Kengy Barty for useful discussions on earlier versions of this paper.

References

[1] Chen, Z. L., W. B. Powell. 1999. A convergent cutting-plane and partial-sampling algorithm for multistage linear programs with recourse. *Journal of Optimization Theory and Applications* 497–524.

[2] Donohue, C. J. 1996. Stochastic network programming and the dynamic vehicle allocation problem. PhD dissertation, University of Michigan, Ann Arbor, Michigan.

[3] Donohue, C. J., J. R. Birge. 2006. The abridged nested decomposition method for multistage stochastic linear programs with relatively complete recourse. *Algorithmic Operations Research* **1** 20–30.

[4] Finardi, E.C., E.L. da Silva. 2006. Solving the hydro unit commitment problem via dual decomposition and sequential quadratic programming. *IEE Transactions on Power Systems* **21**(2) 835–844.

[5] Grimmett, G. R., D. R. Stirzaker. 1992. *Probability and Random Processes*. 2nd ed. Clarendon Press, Oxford.

[6] Hindsberger, M., A. B. Philpott. 2001. ReSa: A method for solving multi-stage stochastic linear programs. SPIX Stochastic Programming Symposium, Berlin.

[7] Linowsky, K., A. B. Philpott. 2005. On the convergence of sampling-based decomposition algorithms for multistage stochastic programs. *Journal of Optimization Theory and Applications* **125** 349–366.

[8] Pereira, M. V. F., L. M. V. G. Pinto. 1991. Multi-stage stochastic optimization applied to energy planning. *Mathematical Programming* **52**(2) 359–375. doi:http://dx.doi.org/10.1007/BF01582895.

[9] Philpott, A. B., Z. Guan. 2008. On the convergence of stochastic dual dynamic programming and related methods. *Operations Research Letters* **36**(4) 450–455.

[10] Philpott, A. B., Z. Guan. 2011. A multistage stochastic programming model for the New Zealand dairy industry. *International Journal of Production Economics* **134**(2) 289–299.

[11] Ruszczynski, A. 2006. *Nonlinear Optimization*. Princeton University Press, Princeton, NJ, USA.

[12] Shapiro, A. 2011. Analysis of Stochastic Dual Dynamic Programming Method. *European Journal of Operational Research* **209** 63–72.