

Exact and converging bounds for SDDP

V. Leclère, P. Carpentier, J-P. Chancelier, A. Lenoir, F. Pacaud

CMS 2018 – NTNU – Trondheim
29/05/2018

Contents

- 1 Introduction
 - Setting
 - Strength and weaknesses of SDDP
- 2 Abstract SDDP
 - Linear Bellman Operator
 - Abstract SDDP
- 3 Dual SDDP
 - Fenchel transform of LBO
 - Dual SDDP
 - Converging upper bound and stopping test
 - Inner Approximation
- 4 Numerical results

Introduction

We are interested in multistage stochastic optimization problems of the form

$$\begin{aligned} \min_{\pi} \quad & \mathbb{E} \left(\sum_{t=0}^{T-1} L_t(\mathbf{X}_t, \mathbf{U}_t, \boldsymbol{\xi}_t) + K(\mathbf{X}_T) \right) \\ \text{s.t.} \quad & \mathbf{X}_{t+1} = f_t(\mathbf{X}_t, \mathbf{U}_t, \boldsymbol{\xi}_t) \\ & \mathbf{U}_t = \pi_t(\mathbf{X}_t, \boldsymbol{\xi}_t) \end{aligned}$$

where

- \mathbf{x}_t is the **state** of the system,
- \mathbf{u}_t is the **control** applied at time t ,
- $\boldsymbol{\xi}_t$ is the **noise** happening between time t and $t + 1$, assumed to be time-independent,
- π is the **policy**.

Stochastic Dynamic Programming

By the white noise assumption, this problem can be solved by **Dynamic Programming**, where the Bellman functions satisfy

$$\begin{cases} V_T(x) &= K(x) \\ \hat{V}_t(x, \xi) &= \min_{u_t \in \mathbb{U}} L_t(x, u_t, \xi) + V_{t+1} \circ f_t(x, u_t, \xi) \\ V_t(x) &= \mathbb{E} \left(\hat{V}_t(x, \xi_t) \right) \end{cases}$$

Indeed, π is an optimal policy if

$$\pi_t(x, \xi) \in \arg \min_{u_t \in \mathbb{U}} \{ L_t(x, u_t, \xi) + V_{t+1} \circ f_t(x, u_t, \xi) \}$$

Bellman operator

For any time t , and any function $R : \mathbb{X} \rightarrow \mathbb{R} \cup \{+\infty\}$ we define

$$\hat{\mathcal{T}}_t(R)(x, \xi) := \min_{u_t \in \mathbb{U}} L_t(x, u_t, \xi) + R \circ f_t(x, u_t, \xi)$$

and

$$\mathcal{T}_t(R)(x) := \mathbb{E} \left[\hat{\mathcal{T}}_t(R)(x, \xi) \right].$$

Thus the Bellman equation simply reads

$$\begin{cases} V_T &= K \\ V_t &= \mathcal{T}_t(V_{t+1}) \end{cases}$$

Incidentally, R induce a policy $\pi_t^R(x, \xi)$

Bellman operator

For any time t , and any function $R : \mathbb{X} \rightarrow \mathbb{R} \cup \{+\infty\}$ we define

$$\hat{\mathcal{T}}_t(R)(x, \xi) := \min_{u_t \in \mathbb{U}} L_t(x, u_t, \xi) + R \circ f_t(x, u_t, \xi)$$

and

$$\mathcal{T}_t(R)(x) := \mathbb{E} \left[\hat{\mathcal{T}}_t(R)(x, \xi) \right].$$

Thus the Bellman equation simply reads

$$\begin{cases} V_T &= K \\ V_t &= \mathcal{T}_t(V_{t+1}) \end{cases}$$

Incidentally, R induce a policy $\pi_t^R(x, \xi)$

Bellman operator

For any time t , and any function $R : \mathbb{X} \rightarrow \mathbb{R} \cup \{+\infty\}$ we define

$$\hat{\mathcal{T}}_t(R)(x, \xi) := \min_{u_t \in \mathbb{U}} L_t(x, u_t, \xi) + R \circ f_t(x, u_t, \xi)$$

and

$$\mathcal{T}_t(R)(x) := \mathbb{E} \left[\hat{\mathcal{T}}_t(R)(x, \xi) \right].$$

Thus the Bellman equation simply reads

$$\begin{cases} V_T &= K \\ V_t &= \mathcal{T}_t(V_{t+1}) \end{cases}$$

Incidentally, R induce a policy $\pi_t^R(x, \xi)$

SDDP algorithm

Under linear dynamics, and convex costs, the SDDP algorithm iteratively constructs polyhedral outer approximations of V_t .

More precisely, at iteration k

- We have polyhedral functions $\underline{V}_t^k(\cdot) = \max_{\kappa \leq k} \langle \lambda_t^\kappa, \cdot \rangle + \beta_t^\kappa$, such that $\underline{V}_t^k \leq V_t$.
- **Forward pass:** We simulate the dynamical system, along one scenario, according to policy $\pi^{\underline{V}_t^k}$, yielding a trajectory $\{\underline{x}_t^k\}_{t \in \llbracket 0, T \rrbracket}$.
- **Backward pass:** We compute cuts $x \mapsto \langle \lambda_t^{k+1}, \cdot \rangle + \beta_t^{k+1} \leq V_t$ along this trajectory, and update our outer approximations.

Contents

- 1 Introduction
 - Setting
 - Strength and weaknesses of SDDP
- 2 Abstract SDDP
 - Linear Bellman Operator
 - Abstract SDDP
- 3 Dual SDDP
 - Fenchel transform of LBO
 - Dual SDDP
 - Converging upper bound and stopping test
 - Inner Approximation
- 4 Numerical results

SDDP strengths

- SDDP is a widely used algorithm in the energy community, with multiple **applications** in
 - mid and long term water storage management problem,
 - long-term investment problems,
 - ...
- Recent works have presented **extensions** of the algorithm to
 - deal with some non-convexity,
 - treat risk-averse or distributionally robust problems,
 - incorporate integer variables.
- Multiple **numerical improvements** have been proposed
 - cut selection
 - regularization
 - multi-cut or ϵ -resolution

SDDP weaknesses

There are still some gaps in our knowledge of this approach:

- there is no convergence speed guaranteed,
- regularization methods are not mature yet,
- there is no good **stopping test**.

SDDP Stopping test

- Exact lower bound of the problem : $\underline{V}_0^k(x_0)$.
- Upper-bound estimated by Monte-Carlo simulation yielding costly statistical stopping tests (Pereira Pinto (1991) or Shapiro (2011))
- Alternative statistical tests have been proposed (see Homem de Mello et al (2011))
- Exact upper-bound computation has been proposed by Philpott et al (2013) but without any proof of convergence, leading to possibly not converging stopping tests.

Linear Bellman Operator

An operator $\mathcal{B} : F(\mathbb{R}^{n_x}) \rightarrow F(\mathbb{R}^{n_x})$ is said to be a *linear Bellman operator* (LBO) if it is defined as follows

$$\mathcal{B}(R) : x \mapsto \inf_{(\mathbf{u}, \mathbf{y})} \mathbb{E} \left[\mathbf{c}^\top \mathbf{u} + R(\mathbf{y}) \right]$$
$$\text{s.t. } T\mathbf{x} + \mathcal{W}_u(\mathbf{u}) + \mathcal{W}_y(\mathbf{y}) \leq \mathbf{h}$$

where $\mathcal{W}_u : \mathcal{L}^0(\mathbb{R}^{n_u}) \rightarrow \mathcal{L}^0(\mathbb{R}^{n_c})$ and $\mathcal{W}_y : \mathcal{L}^0(\mathbb{R}^{n_x}) \rightarrow \mathcal{L}^0(\mathbb{R}^{n_c})$ are two **linear** operators. We denote $S(R)(x)$ the set of \mathbf{y} that are part of optimal solutions to the above problem.

We also define $\mathcal{G}(x)$

$$\mathcal{G}(x) := \{(\mathbf{u}, \mathbf{y}) \mid T\mathbf{x} + \mathcal{W}_u(\mathbf{u}) + \mathcal{W}_y(\mathbf{y}) \leq \mathbf{h}\} .$$

Examples

- Linear point-wise operator:

$$\begin{aligned} \mathcal{W} : \quad \mathcal{L}^0(\mathbb{R}^{n_x}) &\rightarrow \mathcal{L}^0(\mathbb{R}^{n_c}) \\ (\omega \mapsto \mathbf{y}(\omega)) &\mapsto (\omega \mapsto \mathbf{A}\mathbf{y}(\omega)) \end{aligned}$$

Such an operator allows to encode **almost sure constraints**.

- Linear expected operator:

$$\begin{aligned} \mathcal{W} : \quad \mathcal{L}^0(\mathbb{R}^{n_x}) &\rightarrow \mathcal{L}^0(\mathbb{R}^{n_c}) \\ (\omega \mapsto \mathbf{y}(\omega)) &\mapsto (\omega \mapsto \mathbf{A}\mathbb{E}(\mathbf{y})) \end{aligned}$$

Such an operator allows to encode **constraints in expectation**.

Relatively Complete Recourse and cuts

Definition (Relatively Complete Recourse)

We say that the pair (\mathcal{B}, R) satisfy a *relatively complete recourse* (RCR) assumption if for all $x \in \text{dom}(\mathcal{G})$ there exists admissible controls $(\mathbf{u}, \mathbf{y}) \in \mathcal{G}(x)$ such that $\mathbf{y} \in \text{dom}(R)$.

Cut

If R is proper and polyhedral, with RCR assumption, then $\mathcal{B}(R)$ is a proper polyhedral function.

Furthermore, computing $\mathcal{B}(R)(x)$ consists of solving a linear problem which also generates a supporting hyperplane of $\mathcal{B}(R)$, that is, a pair $(\lambda, \beta) \in \mathbb{R}^{n_x} \times \mathbb{R}$ such that

$$\begin{cases} \langle \lambda, \cdot \rangle + \beta \leq \mathcal{B}(R)(\cdot) \\ \langle \lambda, x \rangle + \beta = \mathcal{B}(R)(x) . \end{cases}$$

Setting

Consider a *compatible* sequence of LBO $\{\mathcal{B}_t\}_{t \in \llbracket 0, T-1 \rrbracket}$, that is, such that all admissible controls of \mathcal{B}_t lead to admissible states of \mathcal{B}_{t+1} .

Consider a sequence of functions such that

$$\begin{cases} R_T = K \\ R_t = \mathcal{B}_t(R_{t+1}) \quad \forall t \in \llbracket 0, T-1 \rrbracket \end{cases}$$

Then, the abstract SDDP algorithm generates a sequence of lower polyhedral approximations of R_t . In a forward pass it simulates a trajectory of states, along which the approximation is refined in the backward pass.

Setting

Consider a *compatible* sequence of LBO $\{\mathcal{B}_t\}_{t \in \llbracket 0, T-1 \rrbracket}$, that is, such that all admissible controls of \mathcal{B}_t lead to admissible states of \mathcal{B}_{t+1} .

Consider a sequence of functions such that

$$\begin{cases} R_T = K \\ R_t = \mathcal{B}_t(R_{t+1}) \quad \forall t \in \llbracket 0, T-1 \rrbracket \end{cases}$$

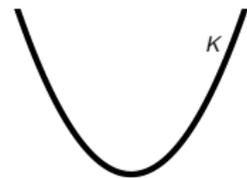
Then, the abstract SDDP algorithm generates a sequence of lower polyhedral approximations of R_t . In a forward pass it simulates a trajectory of states, along which the approximation is refined in the backward pass.

Abstract SDDP

t=0

t=1

t=2



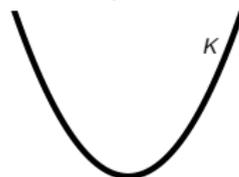
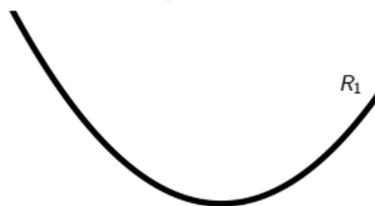
Final Cost $R_2 = K$

Abstract SDDP

t=0

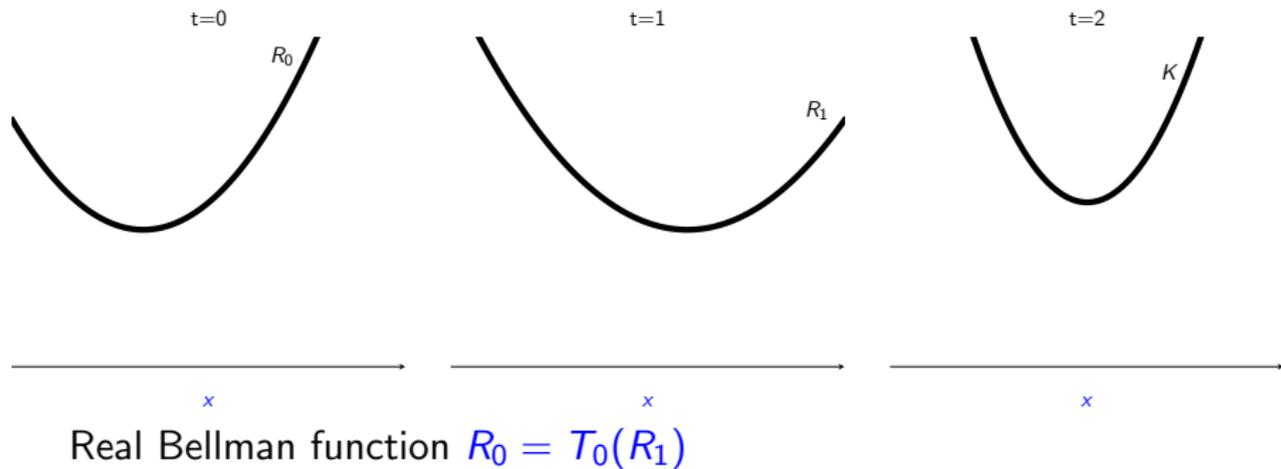
t=1

t=2

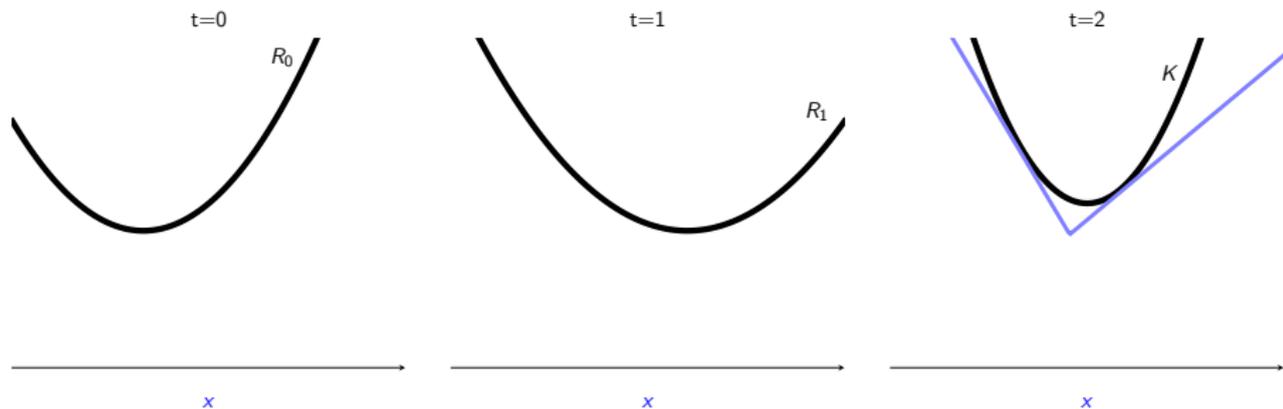


Real Bellman function $R_1 = T_1(R_2)$

Abstract SDDP

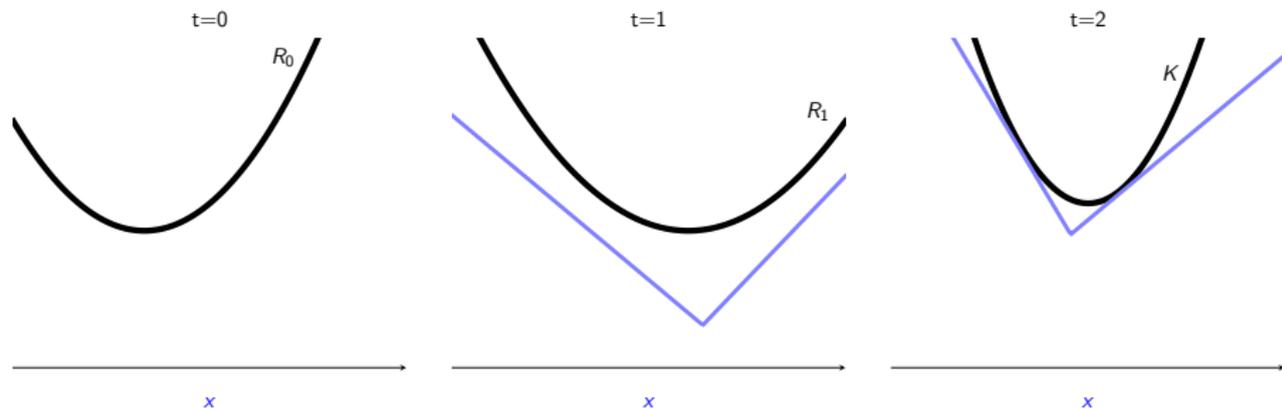


Abstract SDDP



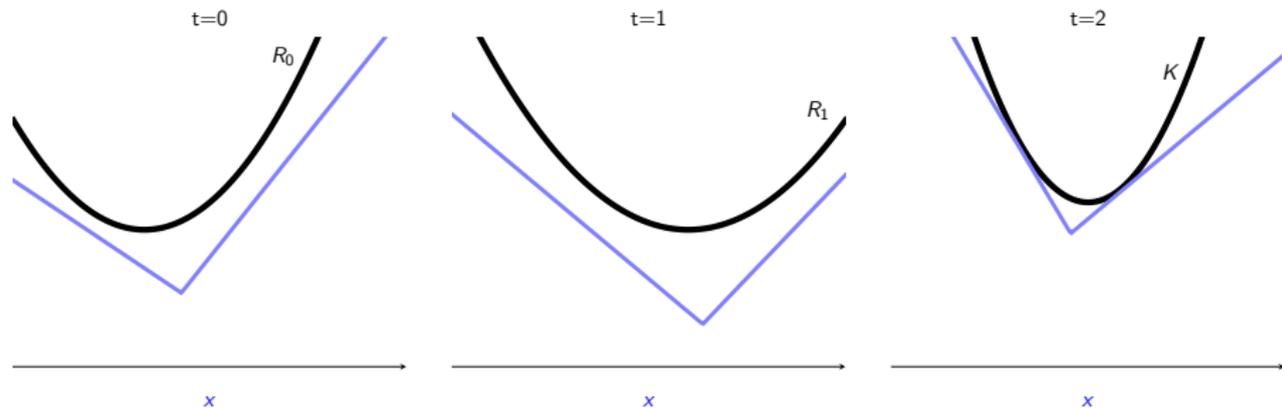
Lower polyhedral approximation \underline{K} of K

Abstract SDDP



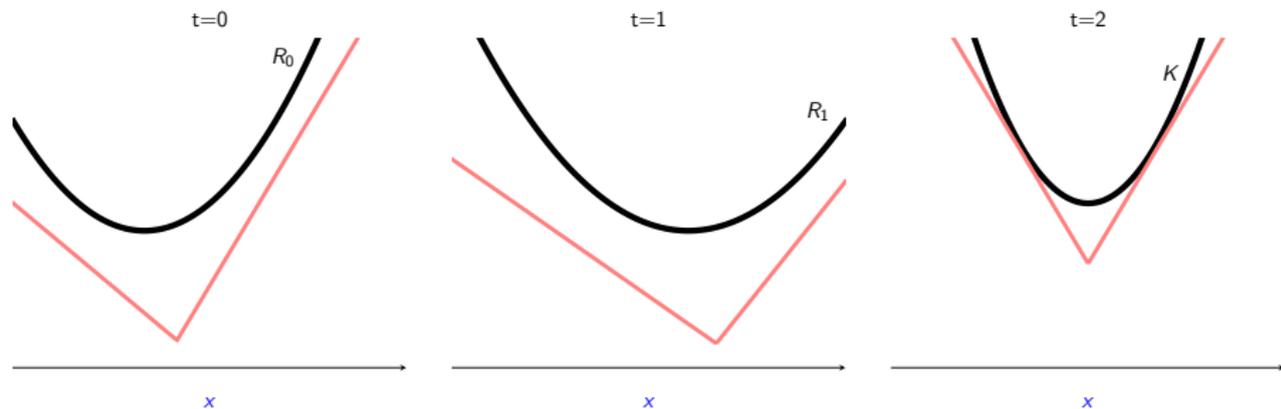
Lower polyhedral approximation $\underline{R}_1 = T_t(\underline{K})$ of R_1

Abstract SDDP



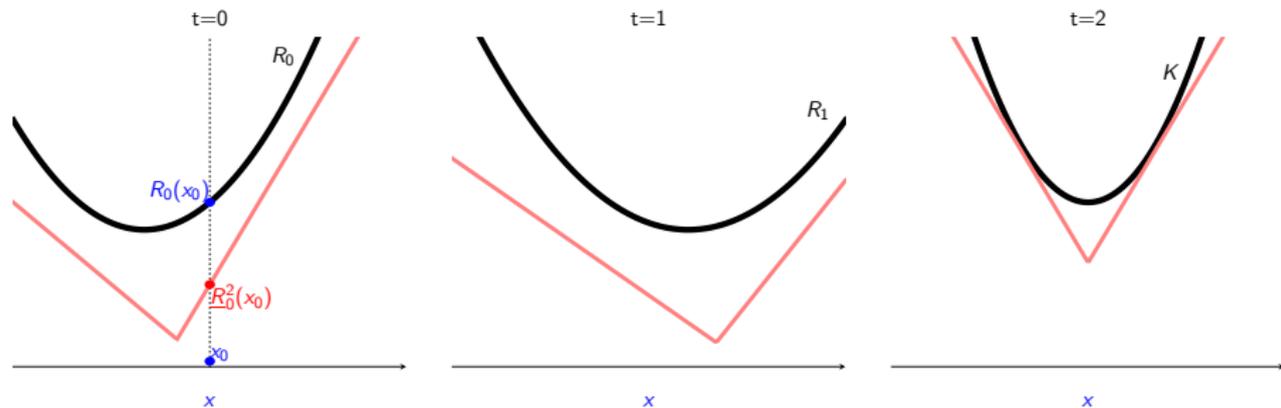
Lower polyhedral approximation $\underline{R}_0 = T_t(\underline{R}_1)$ of R_0

Abstract SDDP



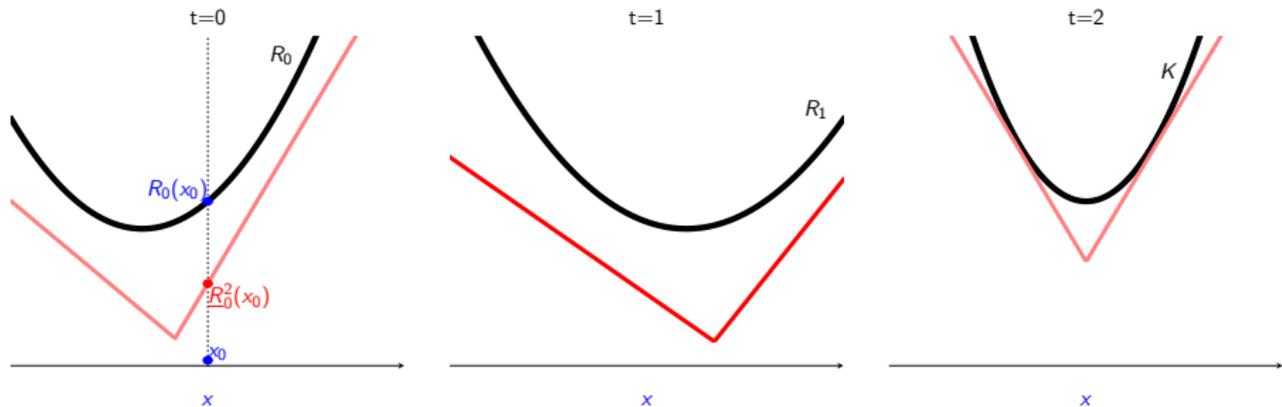
Assume that we have lower polyhedral approximations of R_t

Abstract SDDP



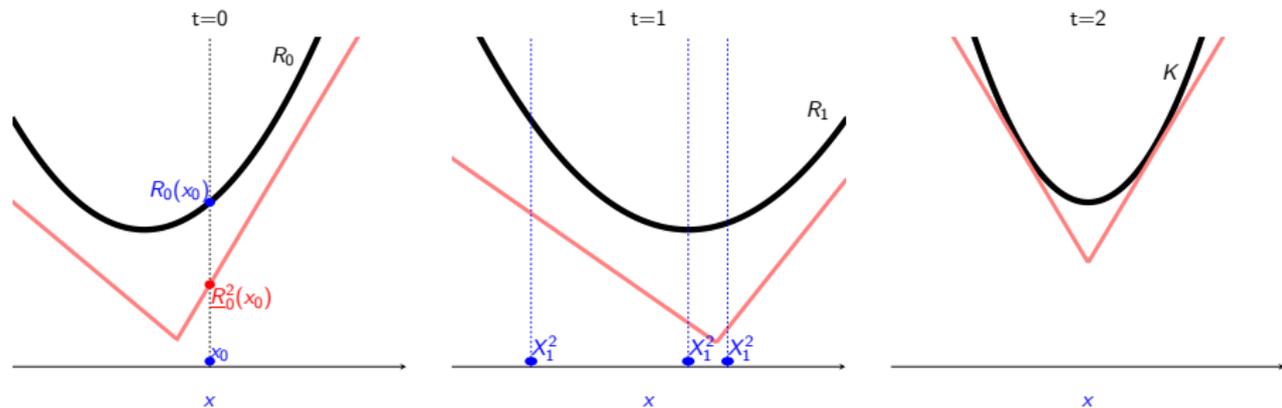
Thus we have a lower bound on the value of our problem

Abstract SDDP



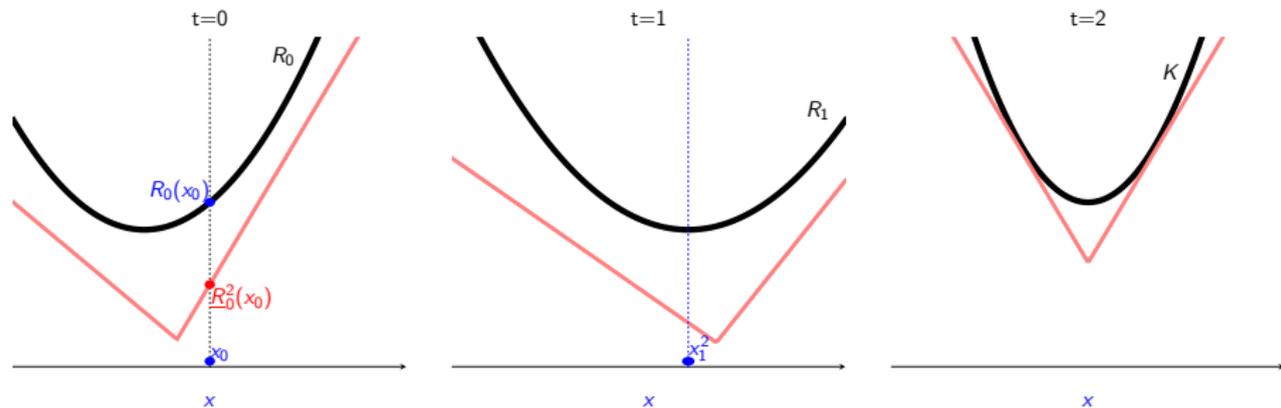
We apply $\pi_0 \frac{R_1^{(2)}}{R_0}$ to x_0 and obtain $x_1^{(2)}$

Abstract SDDP



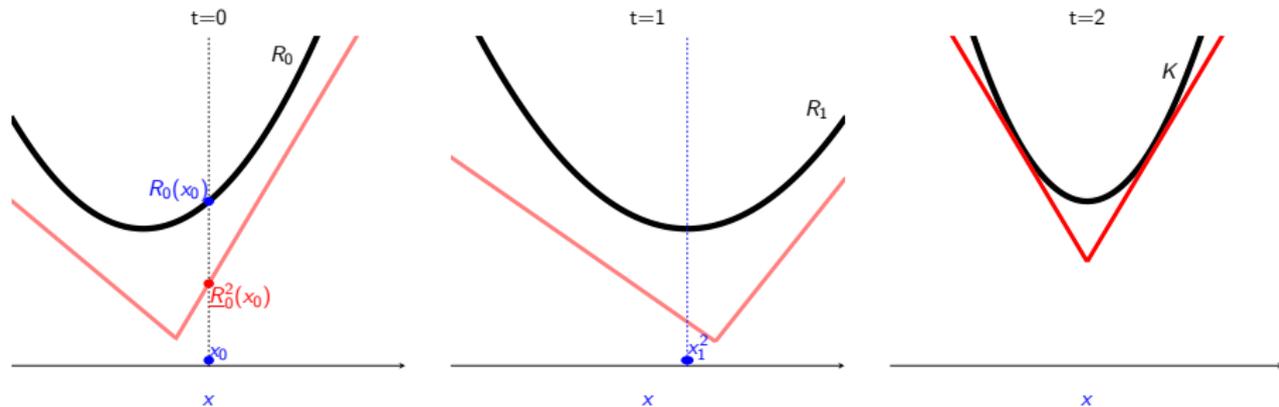
We apply $\pi_0 \frac{R_1^{(2)}}{R_0}$ to x_0 and obtain $x_1^{(2)}$

Abstract SDDP



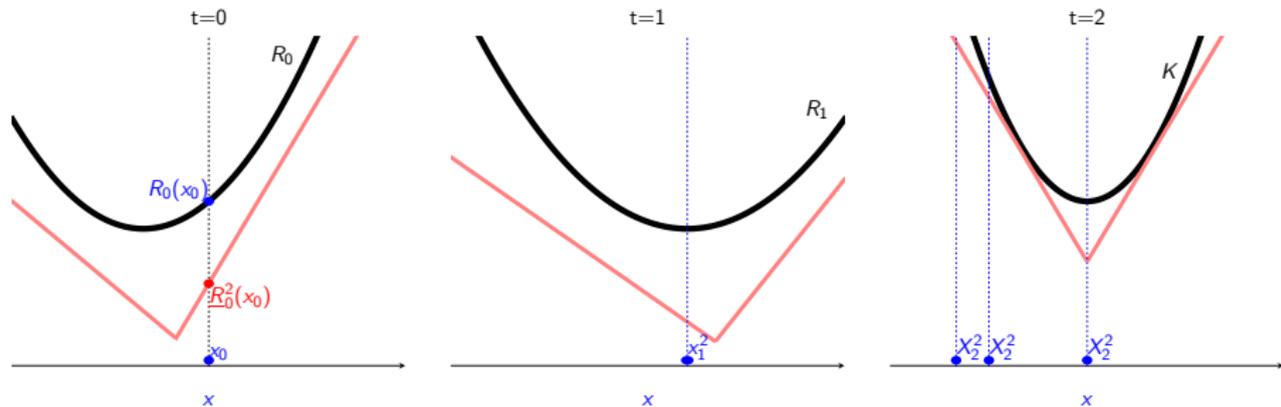
Draw a random realisation $x_1^{(2)}$ of $\mathbf{X}_1^{(2)}$

Abstract SDDP



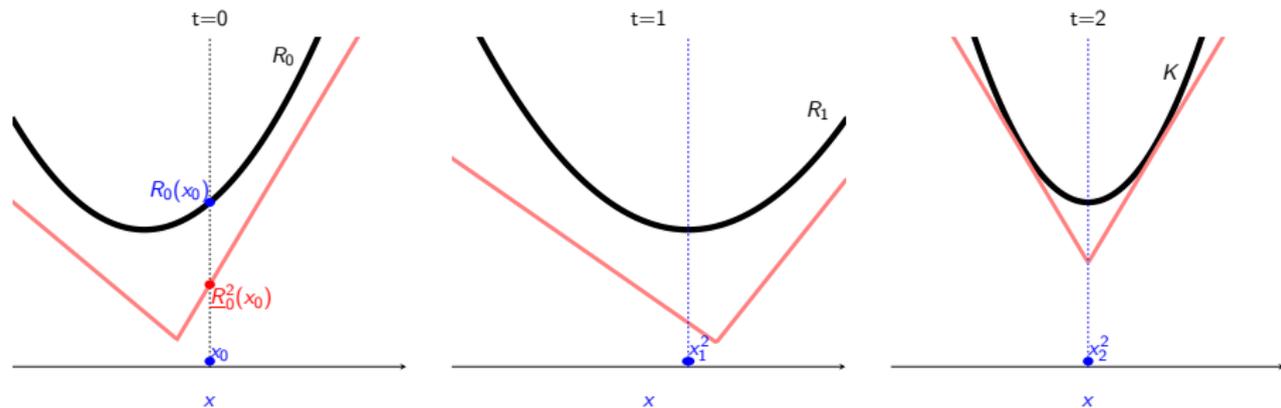
We apply $\pi_1^{R_1^{(2)}}$ to $x_1^{(2)}$ and obtain $x_2^{(2)}$

Abstract SDDP



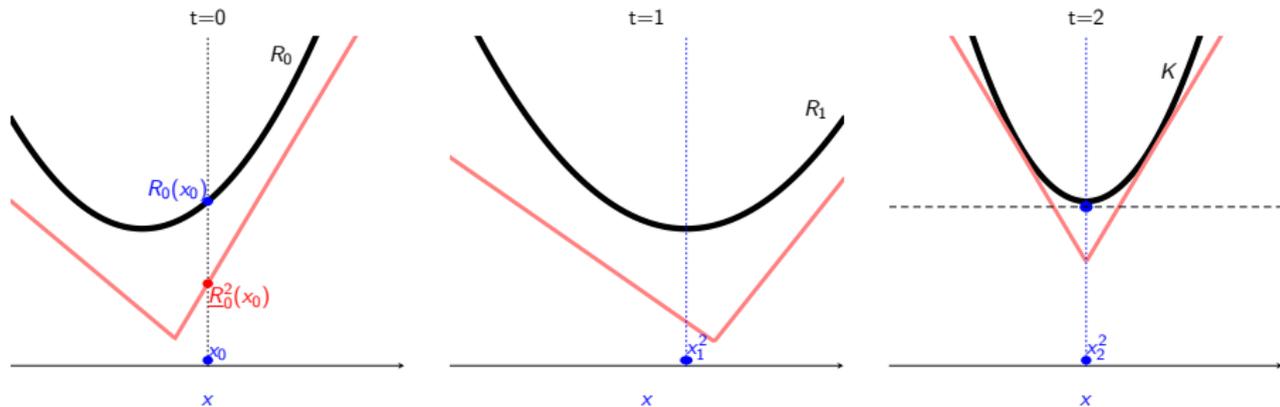
We apply $\pi_1^{R_1^{(2)}}$ to $x_1^{(2)}$ and obtain $x_2^{(2)}$

Abstract SDDP



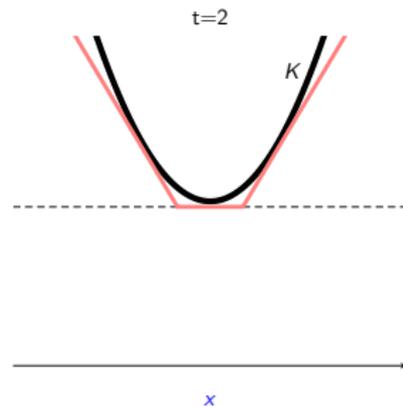
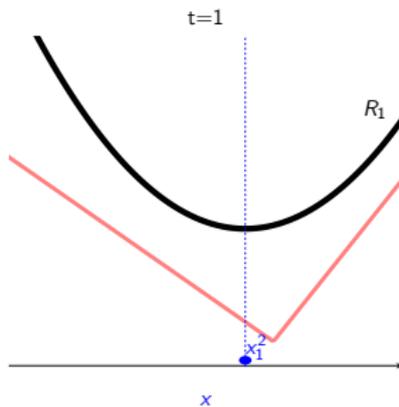
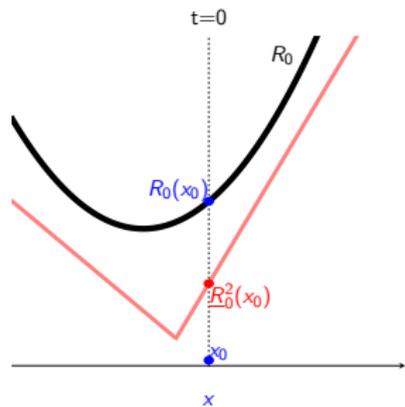
Draw a random realisation $x_2^{(2)}$ of $\mathbf{X}_2^{(2)}$

Abstract SDDP



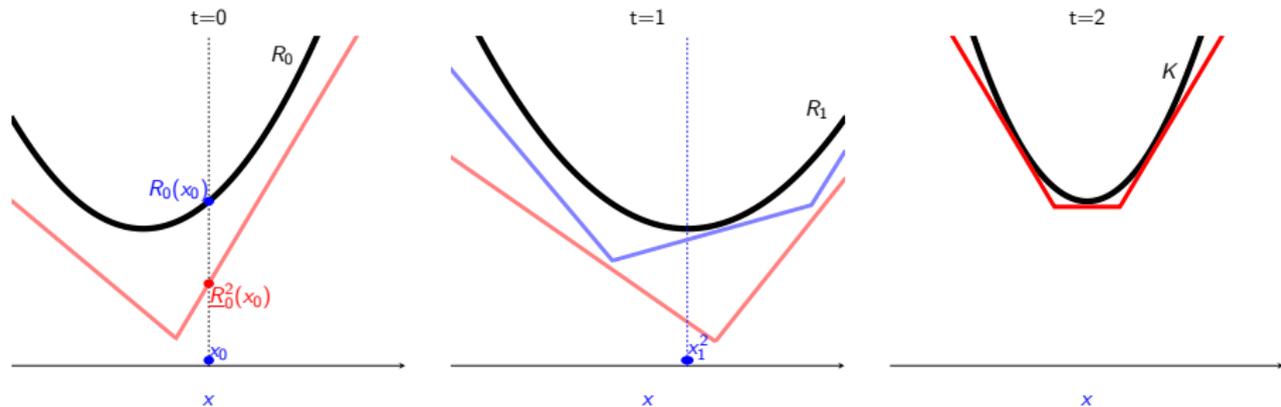
Compute a cut for K at $x_2^{(2)}$

Abstract SDDP



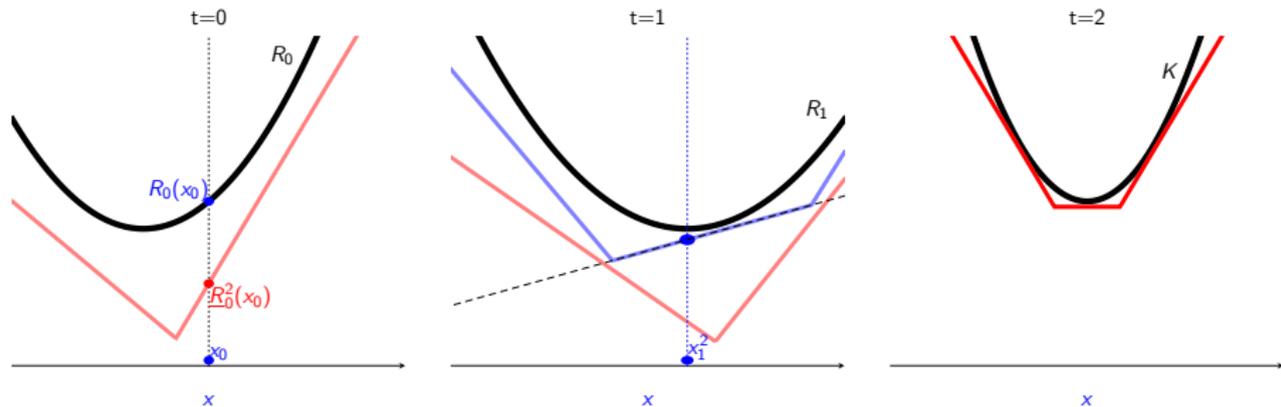
Add the cut to $\underline{R}_2^{(2)}$ which gives $\underline{R}_2^{(3)}$

Abstract SDDP



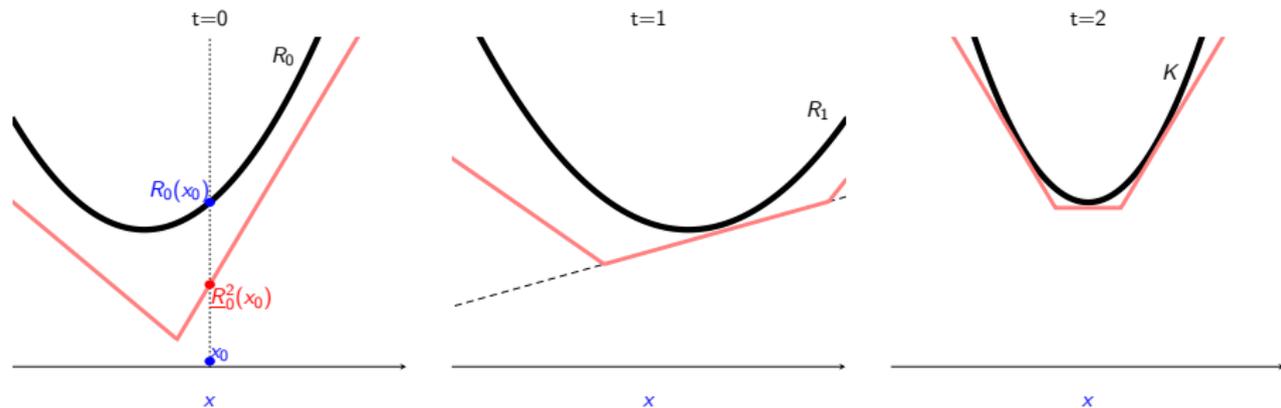
A new lower approximation of R_1 is $T_1(\underline{R}_2^{(3)})$

Abstract SDDP



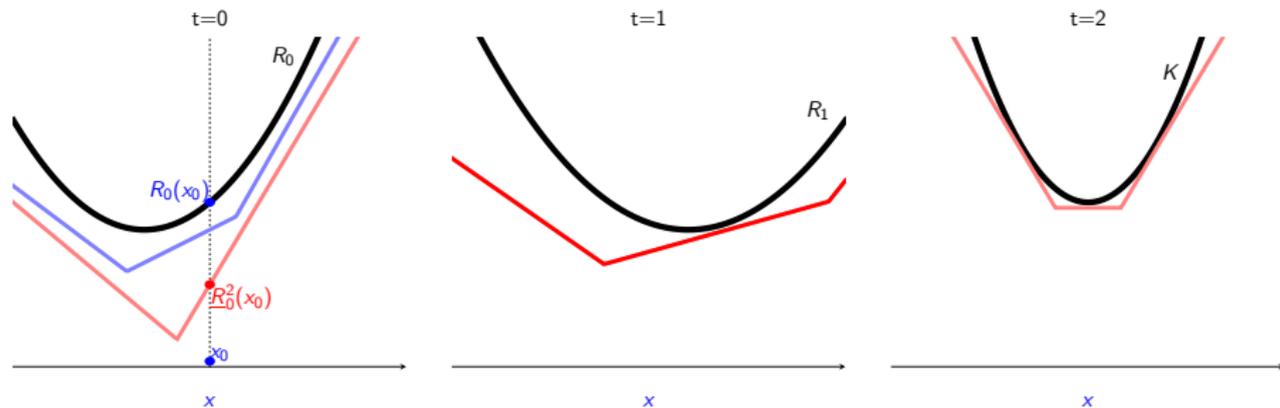
We only compute the face active at $x_1^{(2)}$

Abstract SDDP



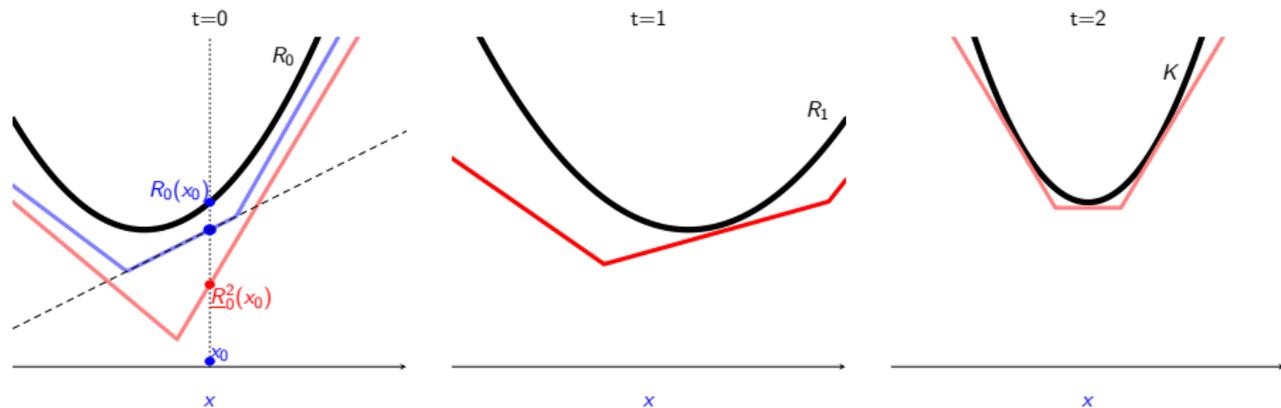
Add the cut to $\underline{R}_1^{(2)}$ which gives $\underline{R}_1^{(3)}$

Abstract SDDP



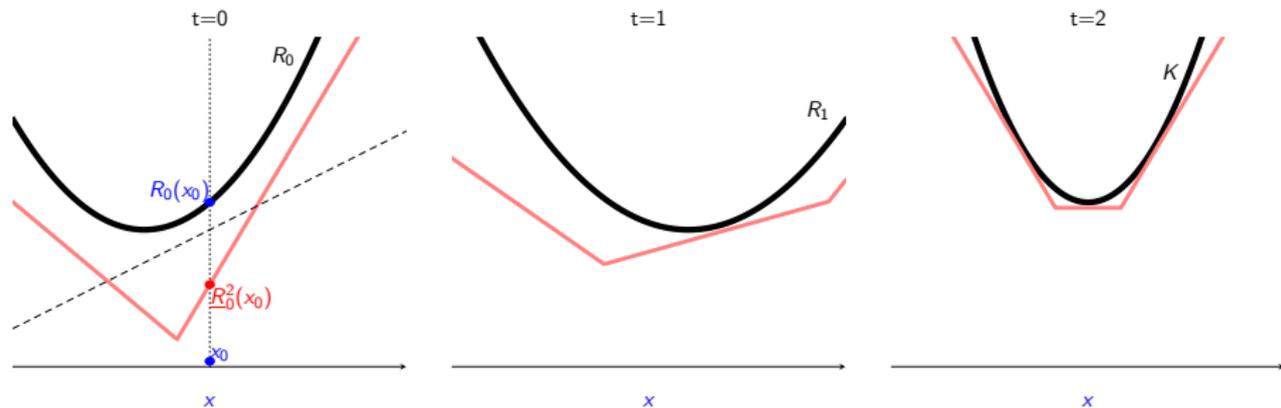
A new lower approximation of R_0 is $T_0(\underline{R}_1^{(3)})$

Abstract SDDP



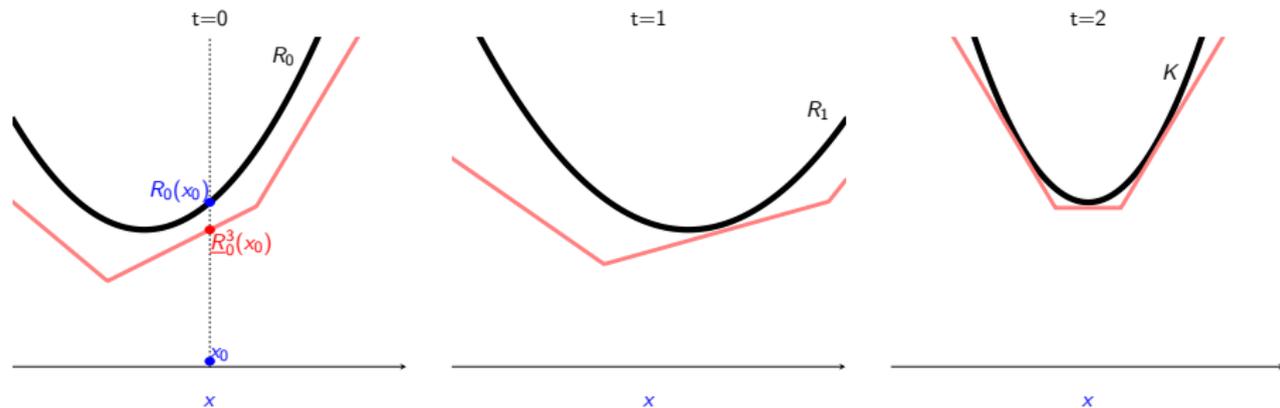
We only compute the face active at x_0

Abstract SDDP



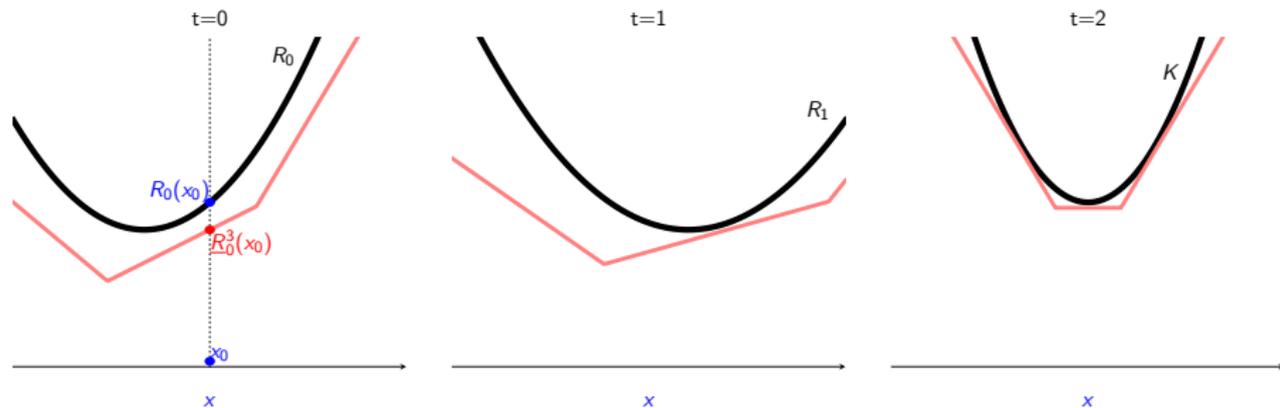
We only compute the face active at x_0

Abstract SDDP



We obtain a new lower bound

Abstract SDDP



We obtain a new lower bound

Abstract SDDP convergence

Theorem

Assume that Ω is finite, $R(x_0)$ is finite, and $\{\mathcal{B}_t\}_t$ is compatible. Further assume that, for all $t \in \llbracket 0, T \rrbracket$ there exists compact sets X_t such that, for all k , $x_t^k \in X_t$ (e.g. \mathcal{B}_t have compact domain).

Then, $(\underline{R}_t^k)_{k \in \mathbb{N}}$ is a non-decreasing sequence of lower approximations of R_t , and $\lim_k \underline{R}_0^k(x_0) = R_0(x_0)$, for $t \in \llbracket 0, T - 1 \rrbracket$.

Further, the cuts coefficients generated remain in a compact set.

Contents

- 1 Introduction
 - Setting
 - Strength and weaknesses of SDDP
- 2 Abstract SDDP
 - Linear Bellman Operator
 - Abstract SDDP
- 3 Dual SDDP
 - Fenchel transform of LBO
 - Dual SDDP
 - Converging upper bound and stopping test
 - Inner Approximation
- 4 Numerical results

Fenchel transform of LBO

Theorem

Assume that the pair (\mathcal{B}, R) satisfy the RCR assumption, R being proper polyhedral, and \mathcal{B} compact (i.e. \mathcal{G} is compact valued with compact domain).

Then $\mathcal{B}(R)$ is a proper function and we have that

$$[\mathcal{B}(R)]^* = \mathcal{B}^\ddagger(R^*)$$

where \mathcal{B}^\ddagger is an explicitly given LBO.

Dual LBO

More precisely we have

$$\begin{aligned} \mathcal{B}^\dagger(Q) : \lambda \mapsto & \inf_{\mu \in \mathcal{L}^0(\mathbb{R}^{n_x}), \nu \in \mathcal{L}^0(\mathbb{R}^{n_c})} \mathbb{E} \left[-\mu^\top \mathbf{h} + Q(\nu) \right] \\ & \text{s.t. } T^\top \mathbb{E}[\mu] + \lambda = 0 \\ & \mathcal{W}_u^\dagger(\mu) = \mathbf{C} \\ & \mathcal{W}_y^\dagger(\mu) = \nu \\ & \mu \leq 0, \end{aligned}$$

Contents

- 1 Introduction
 - Setting
 - Strength and weaknesses of SDDP
- 2 Abstract SDDP
 - Linear Bellman Operator
 - Abstract SDDP
- 3 Dual SDDP
 - Fenchel transform of LBO
 - Dual SDDP
 - Converging upper bound and stopping test
 - Inner Approximation
- 4 Numerical results

Recursion over dual value function

Denote $\mathcal{D}_t := V_t^*$.

Theorem

Then

$$\begin{cases} \mathcal{D}_T &= K^*, \\ \mathcal{D}_t &= \mathcal{B}_t^\dagger(\mathcal{D}_{t+1}) \quad \forall t \in \llbracket 0, T-1 \rrbracket \end{cases}$$

This is a **Bellman recursion** on \mathcal{D}_t instead of V_t .

Recursion over dual value function

Denote $\mathcal{D}_t := V_t^*$.

Theorem

Then

$$\begin{cases} \mathcal{D}_T &= K^*, \\ \mathcal{D}_t &= \mathcal{B}_{t, L_{t+1}}^\dagger(\mathcal{D}_{t+1}) \quad \forall t \in \llbracket 0, T-1 \rrbracket \end{cases}$$

where $\mathcal{B}_{t, L_{t+1}}^\dagger := \mathcal{B}_t^\dagger + \mathbb{I}_{\|\lambda_{t+1}\|_\infty \leq L_{t+1}}$.

This is a **Bellman recursion** on \mathcal{D}_t instead of V_t .

Recursion over dual value function

Denote $\mathcal{D}_t := V_t^*$.

Theorem

Then

$$\begin{cases} \mathcal{D}_T &= K^*, \\ \mathcal{D}_t &= \mathcal{B}_{t, L_{t+1}}^\dagger(\mathcal{D}_{t+1}) \quad \forall t \in \llbracket 0, T-1 \rrbracket \end{cases}$$

where $\mathcal{B}_{t, L_{t+1}}^\dagger := \mathcal{B}_t^\dagger + \mathbb{I}_{\|\lambda_{t+1}\|_\infty \leq L_{t+1}}$.

This is a **Bellman recursion** on \mathcal{D}_t instead of V_t .

Further, under easy technical assumptions, $\{\mathcal{B}_{t, L_{t+1}}^\dagger, t \in \llbracket 0, T \rrbracket\}$ is a compatible sequence of LBOs, where V_t is L_t -Lipschitz.

```

Data: Initial primal point  $x_0$ , Lipschitz bounds  $\{L_t\}_{t \in [0, T]}$ 
for  $k \in \mathbb{N}$  do
  // Forward Pass : compute a set of trial points  $\{\lambda_t^{(k)}\}_{t \in [0, T]}$ 
  Compute  $\lambda_0^k \in \arg \max_{\|\lambda_0\|_\infty \leq L_0} \{x_0^\top \lambda_0 - \underline{D}_0^k(\lambda_0)\}$  ;
  for  $t : 0 \rightarrow T$  do
    select  $\lambda_{t+1}^k \in \arg \min \mathcal{B}_t^\dagger(\underline{D}_{t+1}^k)(\lambda_t^k)$  ;
    and draw a realization  $\lambda_{t+1}^k$  of  $\lambda_{t+1}^k$ ;
  end
  // Backard Pass : refine the lower-approx at trial points
  Set  $\underline{D}_T^k = K^*$  . ;
  for  $t : T - 1 \rightarrow 0$  do
     $\bar{\theta}_t^{k+1} := \mathcal{B}_{t, L_{t+1}}^\dagger(\underline{D}_{t+1}^{k+1})(\lambda_t^k)$  ; // computing cut coefficients
     $\bar{x}_t^{k+1} \in \partial \mathcal{B}_{t, L_{t+1}}^\dagger(\underline{D}_{t+1}^{k+1})(\lambda_t^k)$  ;
     $\bar{\beta}_t^{k+1} := \bar{\theta}_t^{k+1} - \langle \lambda_t^k, \bar{x}_t^{k+1} \rangle$  ;
     $\mathcal{C}_t^{k+1} : \lambda \mapsto \langle \bar{x}_t^{k+1}, \lambda \rangle + \bar{\beta}_t^{k+1}$  ;
     $\underline{D}_t^{k+1} = \max(\underline{D}_t^k, \mathcal{C}_t^{k+1})$  ; // update lower approximation
  end
  If some stopping test is satisfied STOP ;
end
end

```

Contents

- 1 Introduction
 - Setting
 - Strength and weaknesses of SDDP
- 2 Abstract SDDP
 - Linear Bellman Operator
 - Abstract SDDP
- 3 Dual SDDP
 - Fenchel transform of LBO
 - Dual SDDP
 - **Converging upper bound and stopping test**
 - Inner Approximation
- 4 Numerical results

Converging upper bound and stopping test

We have

$$\underline{V}_t^k \leq V_t$$

and

$$\underline{\mathcal{D}}_t^k \leq \mathcal{D}_t \quad \Longrightarrow \quad \underbrace{(\underline{\mathcal{D}}_t^k)^*}_{\approx \bar{V}_t^k} \geq (\mathcal{D}_t^*) = V_t^{**} = V_t$$

Finally, we obtain

$$\underline{V}_0(x_0) \leq V_0(x_0) \leq \bar{V}_0(x_0).$$

Using the convergence of the abstract SDDP algorithm we show that this **bounds are converging**, yielding **converging deterministic stopping tests**.

Converging upper bound and stopping test

We have

$$\underline{V}_t^k \leq V_t$$

and

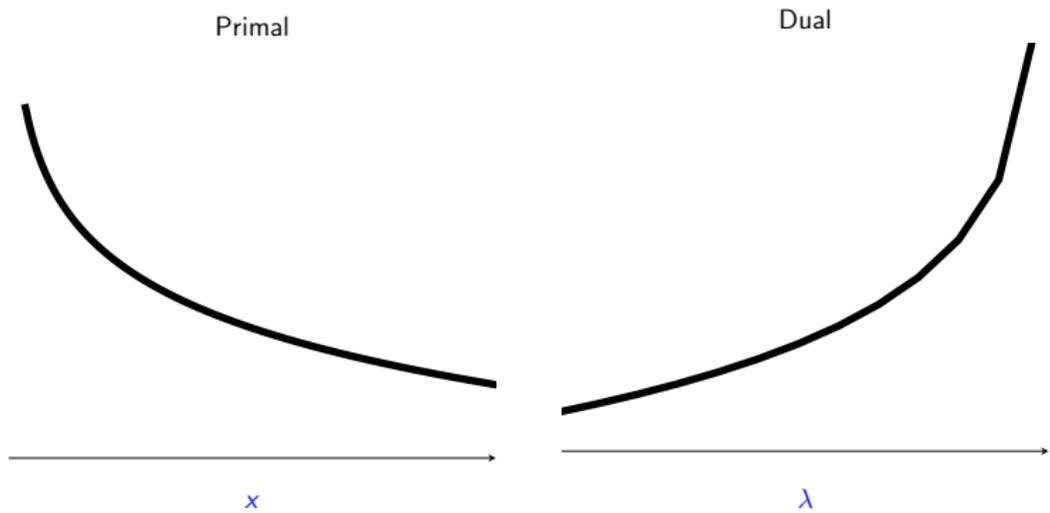
$$\underline{\mathcal{D}}_t^k \leq \mathcal{D}_t \quad \Longrightarrow \quad \underbrace{(\underline{\mathcal{D}}_t^k)^*}_{\approx \bar{V}_t^k} \geq (\mathcal{D}_t^*) = V_t^{**} = V_t$$

Finally, we obtain

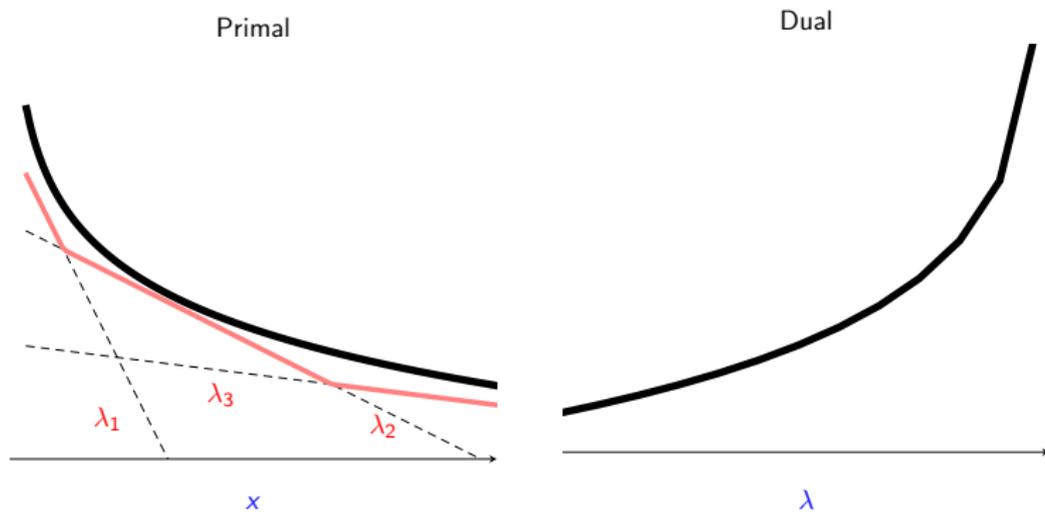
$$\underline{V}_0(x_0) \leq V_0(x_0) \leq \bar{V}_0(x_0).$$

Using the convergence of the abstract SDDP algorithm we show that this **bounds are converging**, yielding **converging deterministic stopping tests**.

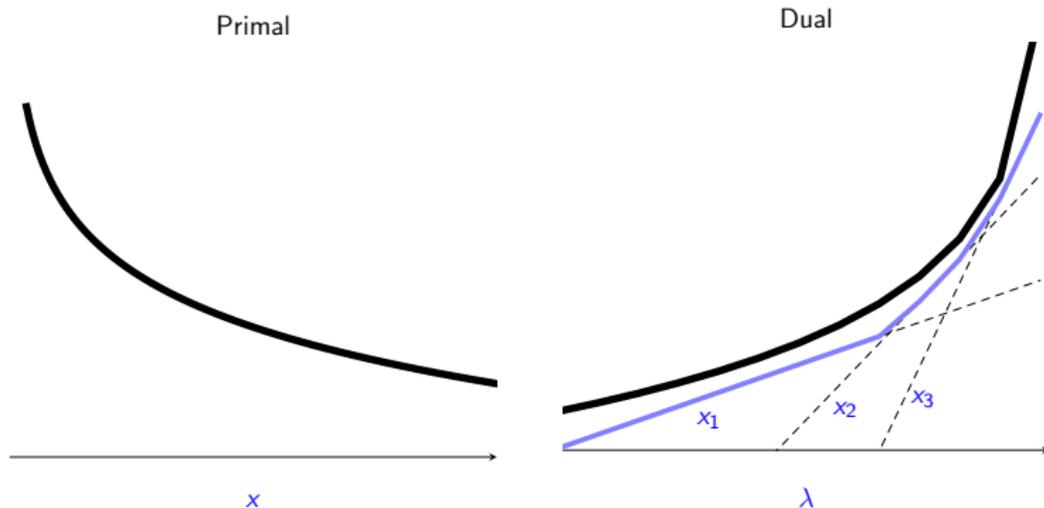
Link between primal and dual approximations



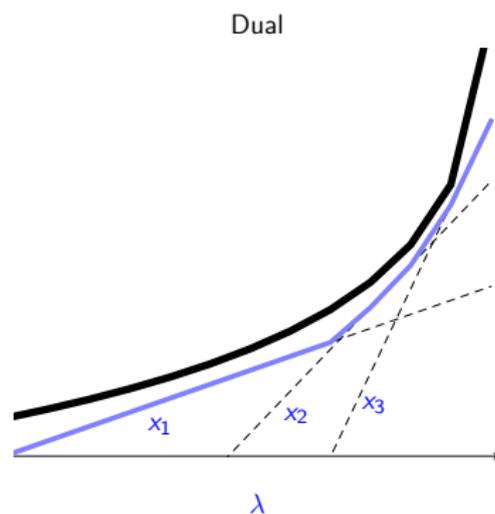
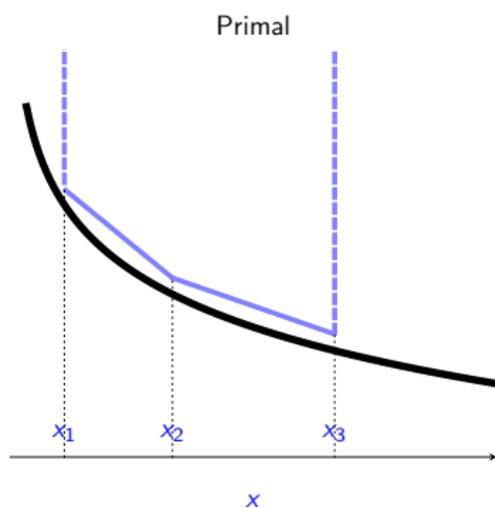
Link between primal and dual approximations



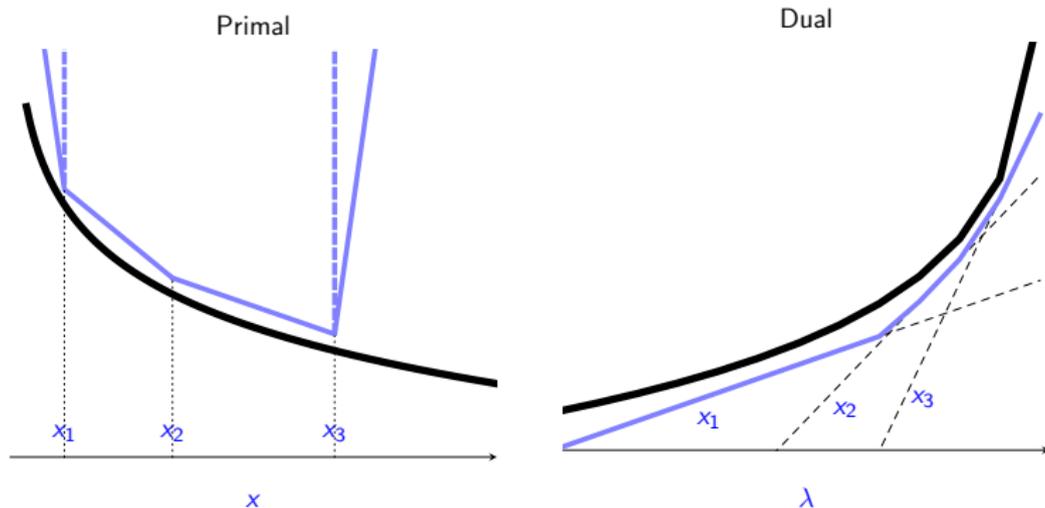
Link between primal and dual approximations



Link between primal and dual approximations



Link between primal and dual approximations



Contents

- 1 Introduction
 - Setting
 - Strength and weaknesses of SDDP
- 2 Abstract SDDP
 - Linear Bellman Operator
 - Abstract SDDP
- 3 Dual SDDP
 - Fenchel transform of LBO
 - Dual SDDP
 - Converging upper bound and stopping test
 - Inner Approximation
- 4 Numerical results

A converging strategy - with guaranteed payoff

Theorem

Let $C_t^{IA,k}(x)$ be the expected cost of the strategy $\pi \bar{V}_t^k$ when starting from state x at time t .

We have,

$$C_t^{IA,k}(x) \leq \bar{V}_t^k(x), \quad \lim_k C_t^{IA,k}(x) = V_t(x)$$

Thus, the inner-approximation yields a new converging strategy, and we have an upper-bound on the (expected) value of this strategy.

Inner Approximation

- $\bar{V}_t^k := [\underline{\mathcal{D}}_t^k]^*$ which is lower than V_t on X_t
- Or

$$\bar{V}_t^k(x) = \min_{\sigma \in \Delta} \left\{ - \sum_{\kappa=1}^k \sigma_{\kappa} \bar{\beta}_t^{\kappa} \quad \left| \quad \sum_{\kappa=1}^k \sigma_{\kappa} \bar{x}_t^{\kappa} = x \right. \right\}$$

- The inner approximation can be computed by solving

$$\begin{aligned} \bar{V}_t^{k+1}(x) &= \sup_{\lambda, \theta} x^{\top} \lambda - \theta \\ \text{s.t.} \quad &\theta \geq \langle \underline{x}_t^i, \lambda \rangle + \bar{\beta}_t^{\kappa} \quad \forall \kappa \in \llbracket 1, k \rrbracket . \end{aligned}$$

Inner Approximation - regularized

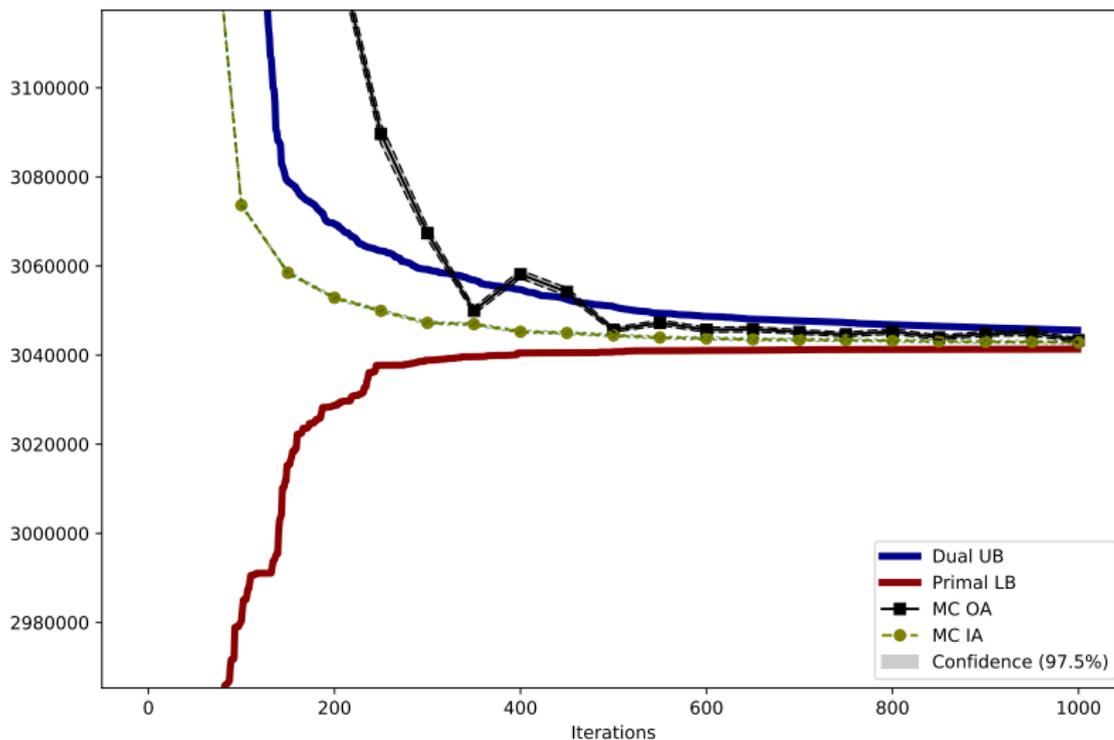
- $\bar{V}_t^k := [\underline{\mathcal{D}}_t^k]^* \square(L_t \|\cdot\|_1)$ which is lower than V_t on X_t
- Or

$$\bar{V}_t^k(x) = \min_{y \in \mathbb{R}^{n_x}, \sigma \in \Delta} \left\{ L_t \|x - y\|_1 - \sum_{\kappa=1}^k \sigma_{\kappa} \bar{\beta}_t^{\kappa} \mid \sum_{\kappa=1}^k \sigma_{\kappa} \bar{x}_t^{\kappa} = y \right\}$$

- The inner approximation can be computed by solving

$$\begin{aligned} \bar{V}_t^{k+1}(x) &= \sup_{\lambda, \theta} x^{\top} \lambda - \theta \\ \text{s.t.} \quad &\theta \geq \langle \underline{x}_t^i, \lambda \rangle + \bar{\beta}_t^{\kappa} \quad \forall \kappa \in \llbracket 1, k \rrbracket . \\ &\|\lambda\|_{\infty} \leq L_t \end{aligned}$$

Numerical results



Stopping test

ε (%)	Dual stopping test		Statistical stopping test	
	n it.	CPU time	n it.	CPU time
2.0	156	183s	250	618s
1.0	236	400s	300	787s
0.5	388	1116s	450	1429s
0.1	> 1000	.	1000	5519s

Table: Comparing dual and statistical stopping criteria for different accuracy levels ε .

Conclusion

- We extend the SDDP algorithm to an abstract framework.
- Leveraging Fenchel conjugate we are able to show a dynamic recursion between dual Bellman value functions.
- We can apply SDDP to this dual recursion.
- This yields a converging exact upper bound on the value of the original problem, hence giving exact and converging stopping tests.
- This also yields a converging strategy with guaranteed payoff.

More information : http://www.optimization-online.org/DB_FILE/2018/04/6575.pdf

Conclusion

- We extend the SDDP algorithm to an abstract framework.
- Leveraging Fenchel conjugate we are able to **show a dynamic recursion between dual Bellman value functions**.
- We can apply SDDP to this dual recursion.
- This yields a **converging exact upper bound** on the value of the original problem, hence giving exact and converging stopping tests.
- This also yields a **converging strategy with guaranteed payoff**.

More information : http://www.optimization-online.org/DB_FILE/2018/04/6575.pdf

Conclusion

- We extend the SDDP algorithm to an abstract framework.
- Leveraging Fenchel conjugate we are able to **show a dynamic recursion between dual Bellman value functions**.
- We can apply SDDP to this dual recursion.
- This yields a **converging exact upper bound** on the value of the original problem, hence giving exact and converging stopping tests.
- This also yields a **converging strategy with guaranteed payoff**.

More information : http://www.optimization-online.org/DB_FILE/2018/04/6575.pdf

Conclusion

- We extend the SDDP algorithm to an abstract framework.
- Leveraging Fenchel conjugate we are able to **show a dynamic recursion between dual Bellman value functions**.
- We can apply SDDP to this dual recursion.
- This yields a **converging exact upper bound** on the value of the original problem, hence giving exact and converging stopping tests.
- This also yields a **converging strategy with guaranteed payoff**.

More information : http://www.optimization-online.org/DB_FILE/2018/04/6575.pdf

Conclusion

- We extend the SDDP algorithm to an abstract framework.
- Leveraging Fenchel conjugate we are able to **show a dynamic recursion between dual Bellman value functions**.
- We can apply SDDP to this dual recursion.
- This yields a **converging exact upper bound** on the value of the original problem, hence giving exact and converging stopping tests.
- This also yields a **converging strategy with guaranteed payoff**.

More information : http://www.optimization-online.org/DB_FILE/2018/04/6575.pdf

Bibliography



Mario VF Pereira and Leontina MVG Pinto.
Multi-stage stochastic optimization applied to energy planning.
Mathematical programming, 52(1-3):359–375, 1991.



Tito Homem-de Mello, Vitor L De Matos, and Erlon C Finardi.
Sampling strategies and stopping criteria for stochastic dual dynamic programming: a case study in long-term hydrothermal scheduling.
Energy Systems, 2(1):1–31, 2011.



Andrew Philpott, Vitor de Matos, and Erlon Finardi.
On solving multistage stochastic programs with coherent risk measures.
Operations Research, 61(4):957–970, 2013.



Alexander Shapiro.
Analysis of stochastic dual dynamic programming method.
European Journal of Operational Research, 209(1):63–72, 2011.



Pierre Girardeau, Vincent Leclere, and Andrew B Philpott.
On the convergence of decomposition methods for multistage stochastic convex programs.
Mathematics of Operations Research, 40(1):130–145, 2014.

Bibliography



Regan Baucke, Anthony Downward, and Golbon Zakeri.

A deterministic algorithm for solving multistage stochastic programming problems.

Optimization Online, 2017.



Vincent Guigues.

Dual dynamic programming with cut selection: Convergence proof and numerical experiments.

European Journal of Operational Research, 258(1):47–57, 2017.



Wim Van Ackooij, Welington de Oliveira, and Yongjia Song.

On regularization with normal solutions in decomposition methods for multistage stochastic programming.

Optimization Online, 2017.