Distributionally Robust SDDP

Andy Philpott[1] Vitor de Matos[2] Lea Kapelevich[1]

[1]Electric Power Optimization Center
University of Auckland

[2]Plan4 Engenharia

EPOC Mini Workshop, 2017

Hydrothermal Scheduling
(in New Zealand)

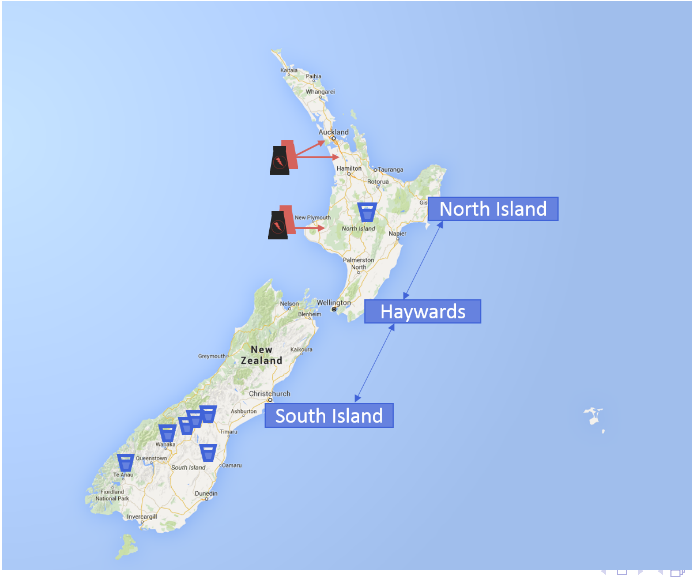

SDDP (Stochastic Dual Dynamic Programming [PP91])

- Policy is defined implicitly through an approximation of a cost-to-go function at every stage;
- This approximation is made of linear cutting planes;
- We add one cut per iteration;
- Each iteration has a *forward pass* and a *backward pass*.

SDDP: Forward Pass

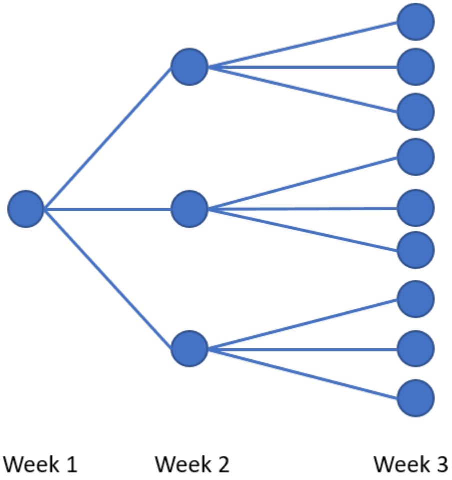

SDDP: Forward Pass

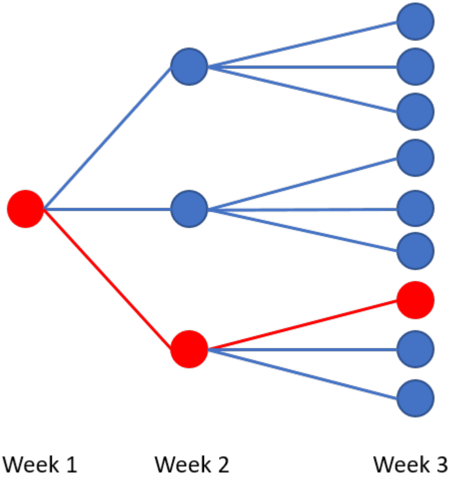

SDDP: Backward Pass

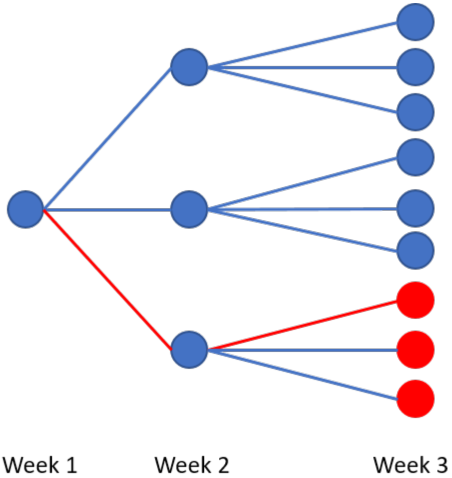

SDDP: Backward Pass

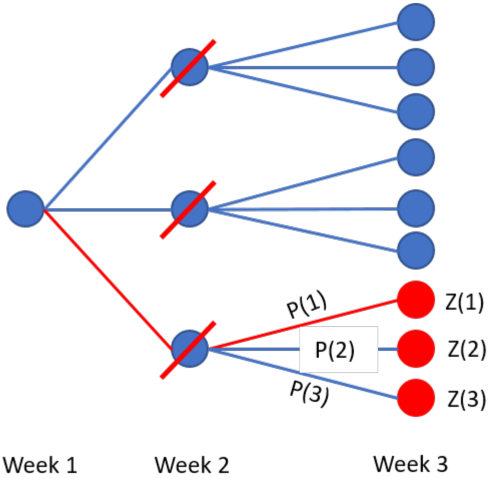

SDDP: Backward Pass

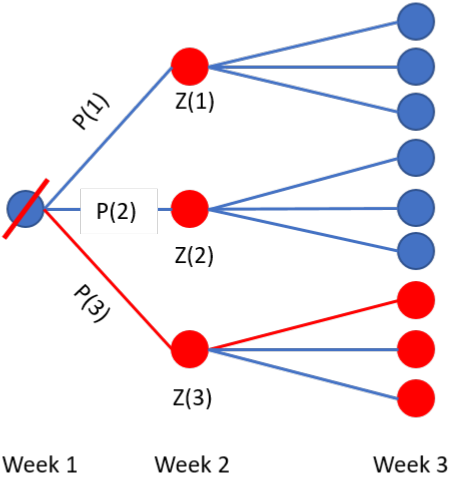

SDDP: Backward Pass

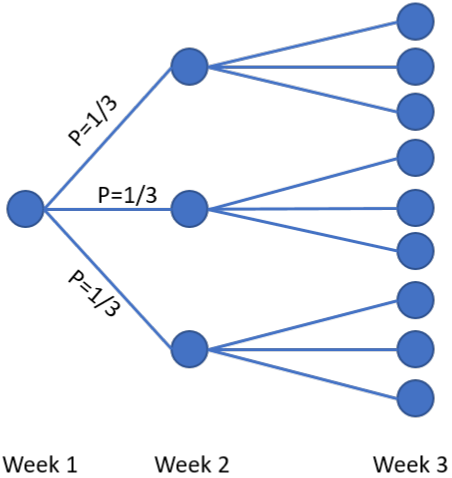

DRO with SDDP

Robust Optimisation

$$\min_{x \in \mathcal{X}} \max_{\omega \in \bar{\Omega}} [Z(x, \omega)].$$

Distributionally Robust Optimisation (DRO)

$$\min_{x \in \mathcal{X}} \max_{\mathbb{P} \in \mathcal{P}} \mathbb{E}_{\mathbb{P}}[Z(x, \omega)].$$

DRO with SDDP

In every stage t,

$$\min_{x_t \in \mathcal{X}_t(x_{t-1}, \omega_t)} \left[c_t^\top x_t + \max_{\mathbb{P} \in \mathcal{P}} \mathbb{E}_{\mathbb{P}}[Z_{t+1}(x_t, \omega_{t+1})] \right].$$

Problem Representation

Need to solve a subproblem of the form:

$$\max_{\mathbb{P}\in\mathcal{P}} \mathbb{E}_{\mathbb{P}}[Z(x,\omega)]$$

where $Z(x,\omega)$ is a cost, and ω is obtained by sampling and takes a finite number of values ω_i, $i = 1, 2, \ldots, m$.
Our choice of $\mathcal{P}$:

$$\mathcal{P} = \{p \in \mathbb{R}^m \mid \sum_{i=1}^{m} p_i = 1,\ p \geq 0,\ \left\| p - \frac{e}{m} \right\|_2 \leq r\}.$$

We can also write the problem as:

$$\begin{array}{lll} \text{P:} & \max & \sum_{i=1}^{m} z_i p_i \\ & \text{s.t.} & \sum_{i=1}^{m} p_i = 1, \\ & & \left\| p - \frac{e}{m} \right\|_2^2 \leq r^2, \\ & & p \geq 0. \end{array} \tag{1}$$

Statistical Interpretation

The convex set $\mathbb{P}$ lies in has a statistical interpretation:

$$\sum_i \frac{(p_i - q_i)^2}{q_i} \leq \chi^2.$$

We use $q_i = \frac{1}{m}$, so q is a vector denoting the uniform distribution with m scenarios.

Geometric Interpretation

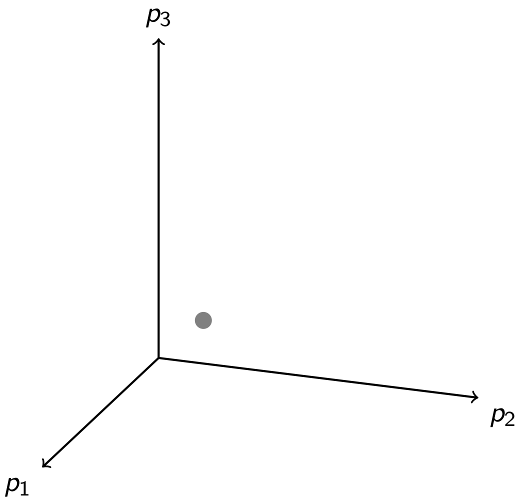

Geometric Interpretation

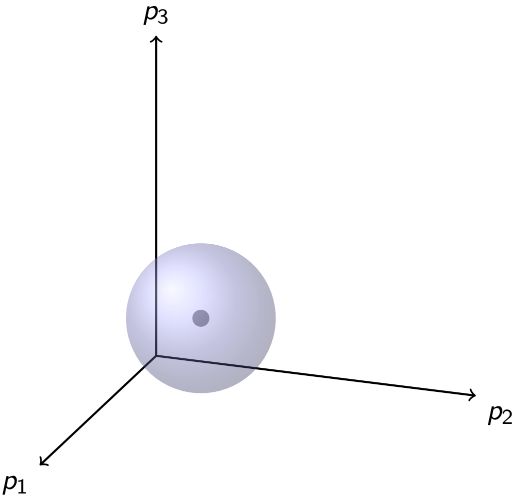

Geometric Interpretation

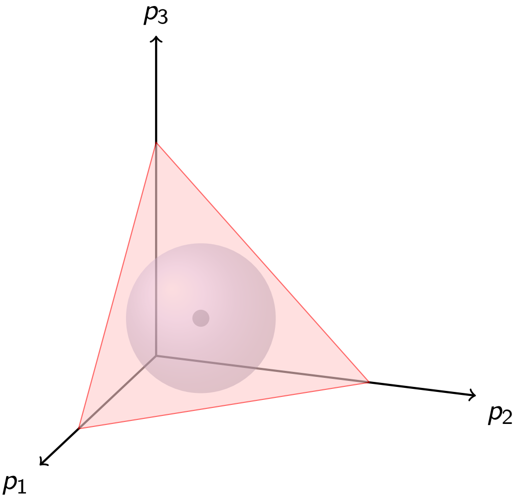

Geometric Interpretation

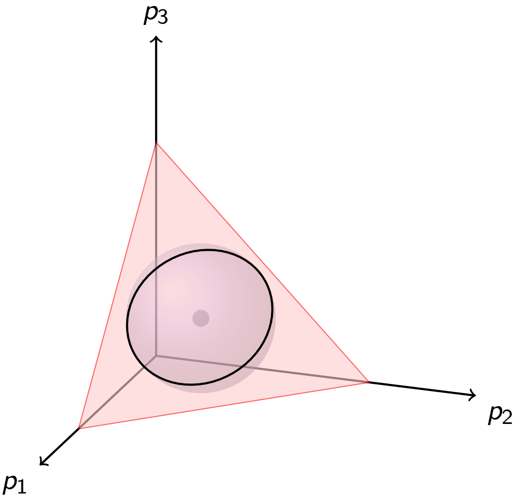

If the radius is small enough, we can drop the $p \geq 0$ constraint. There is a closed form solution for Problem 1:

$$p_i = \frac{1}{m} + \frac{z_i - \bar{z}}{\sqrt{m}} \frac{r}{s}. \tag{2}$$

- $\bar{z}$ is the mean of the future costs z_i,
- s is the standard deviation of future costs z_i.

The objective will be

$$\frac{1}{m} \sum_{i=1}^{m} z_i + \left(\sqrt{m} s\right) r. \tag{3}$$

This is the mean of z_i plus a scaled standard deviation of z_i.

Expanding the Ball

- To drop $p \geq 0$, biggest possible radius is $\frac{1}{\sqrt{S(S-1)}}$ (S scenarios);
- This allows probabilities to change by at most $\frac{\sqrt{S(S-1)}}{S^2}$;
- Not effective when the number of scenarios is moderate, we have $S = 30$;
- But we can't let probabilities become negative.

Larger Radii

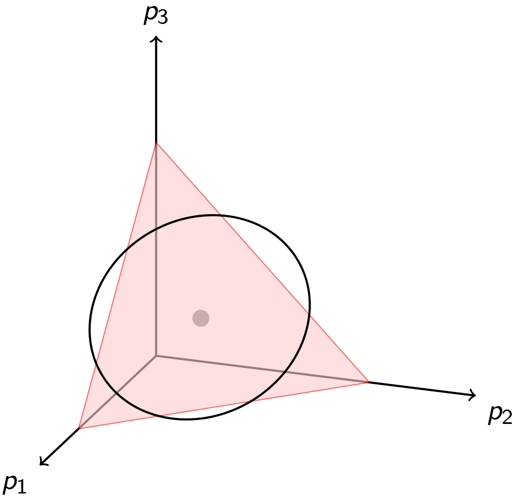

Larger Radii

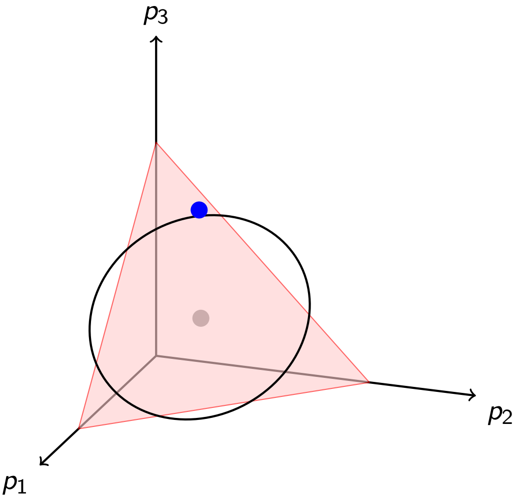

Larger Radii

Fix components of p that are negative to a value of 0, and re-solve in a lower dimension. We have:

$$\begin{array}{rll} \text{P:} & \max & \sum_{i=k+1}^{m} z_i p_i \\ & \text{s.t.} & \sum_{i=k+1}^{m} p_i = 1 - \frac{k}{m}, \\ & & \left\| p - \frac{e}{m} \right\|_2^2 \leq r^2 - \frac{k}{m^2}, \\ & & p \geq 0. \end{array} \tag{4}$$

Final solution looks like:

$$p_i = \begin{cases} 0 & i = 1, \ldots, k, \\ \frac{1}{(m-k)} + \frac{\sqrt{(m-k)r^2 - \frac{k}{m}} \frac{z_i - \bar{z}}{s}}{(m-k)} & i = k+1, \ldots, m. \end{cases} \tag{5}$$

Effectively some scenarios are dropped.

Larger Radii (Example: $r = 2/3$)

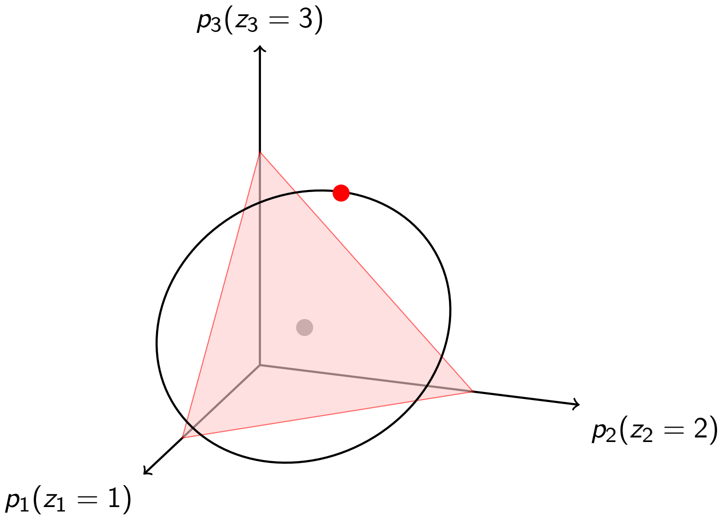

Larger Radii (Example: $r = 2/3$)

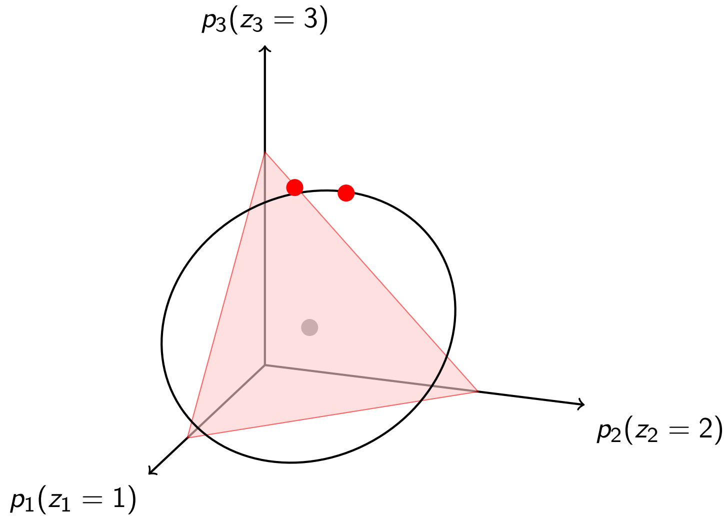

Computational Experiments

- Created policies with and without a DRO approach;
- Each policy consisted of 10,000 cuts;
- Forward passes included inflows from 30 years (1970—1999);
- Simulations were performed out of sample (using inflows from 2000—2009);
- Our technology: Julia [BEKS17], JuMP [DHL17], and SDDP.jl [Dow17].

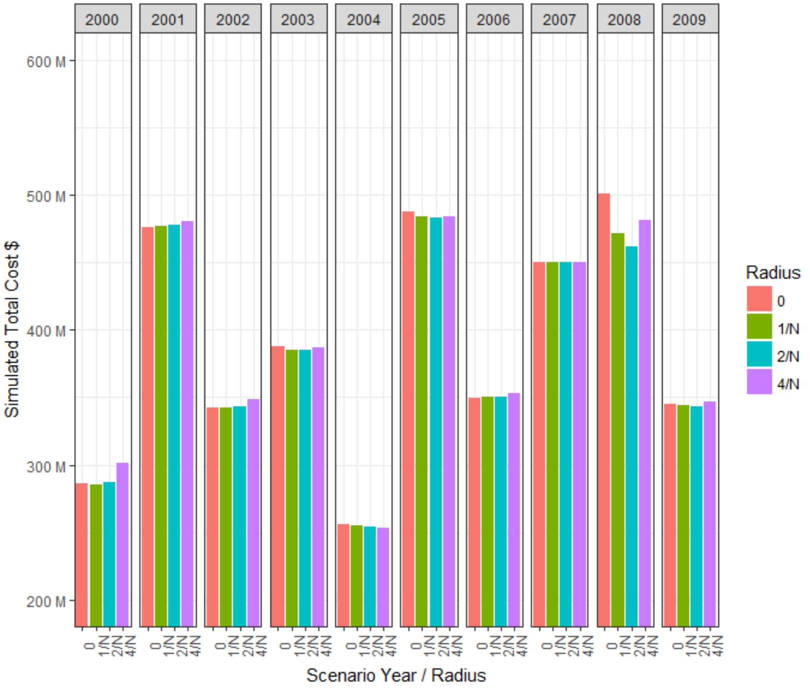
2000
2001
2002
2003
2004
2005
2006
2007
2008
2009
600 M
500 M
400 M
300 M
200 M
Simulated Total Cost $
Radius
0
1/N
2/N
4/N
Scenario Year / Radius

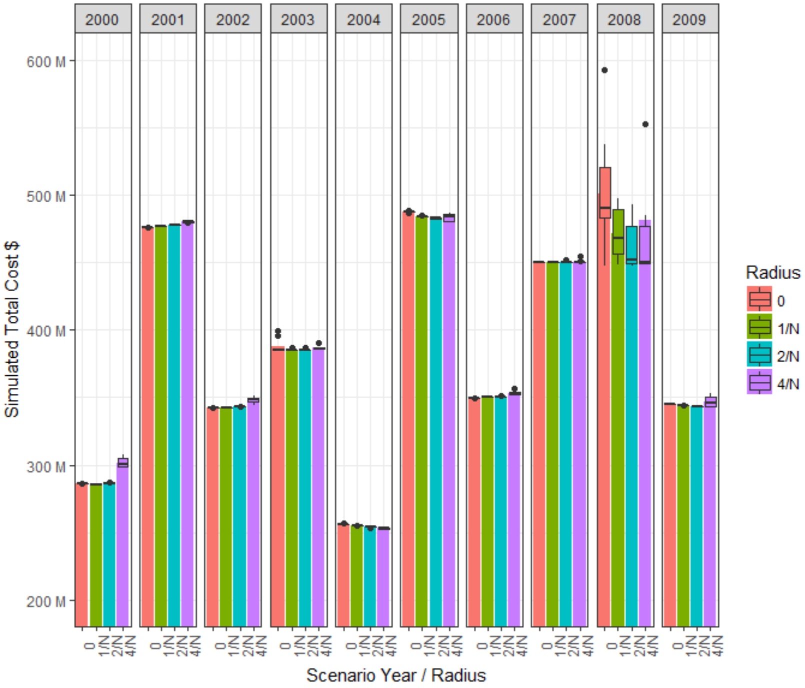

2000
2001
2002
2003
2004
2005
2006
2007
2008
2009
600 M
500 M
400 M
300 M
200 M
Simulated Total Cost $
Scenario Year / Radius
Radius
0
1/N
2/N
4/N

Conclusions

- We can modify the SDDP algorithm to compute policies that are robust against the distribution of an uncertain variable; computing the worst-case distribution is tractable.
- A DRO approach in SDDP allows us to be risk averse without relying on nested risk measures.
- From our experiments, being risk averse in SDDP is desirable not only to reduce costs in years where the risk of water shortage is high, but also to reduce the *variance* within the costs of policies that we compute.

Thank you.

Table: Mean simulated cost ($) for policies created with different uncertainty sets.

Year/Radius	0	1/M	2/M	4/M	AV@F
2000	286,225,815	285,741,825	286,808,544	301,684,926	288,513,
2001	475,933,285	476,990,669	477,778,305	480,111,851	478,570,
2002	342,189,502	342,633,445	343,482,533	348,344,883	344,503,
2003	387,975,202	385,472,658	385,503,156	386,779,738	385,913,
2004	256,211,965	254,957,532	254,188,830	253,308,194	254,152,
2005	487,581,431	484,348,853	482,791,726	483,685,092	481,887,
2006	349,597,459	350,264,052	350,572,763	353,511,872	351,004,
2007	449,998,785	449,802,621	449,987,604	450,603,151	449,802,
2008	501,068,672	471,477,052	461,729,581	481,015,092	466,563,
2009	345,045,595	344,281,225	343,430,918	347,191,191	343,129,

Table: Sample standard deviation ($) for policies created with different uncertainty sets.

Year/Radius	0	1/M	2/M	4/M	AV@R
2000	82,305	137,977	282,351	3,485,192	2,483,391
2001	88,075	39,714	33,128	244,725	260,558
2002	100,076	111,202	331,808	2,090,057	214,437
2003	5,158,708	559,929	367,870	1,295,636	123,560
2004	348,440	140,350	257,759	596,517	466,929
2005	529,447	191,416	188,008	2,720,809	274,674
2006	64,129	203,386	437,754	1,407,594	425,219
2007	8,412	20,964	617,355	1,605,544	31,615
2008	42,335,412	18,742,955	18,423,647	60,107,585	19,668,080
2009	19,058	24,869	26,291	3,964,426	156,931

Results

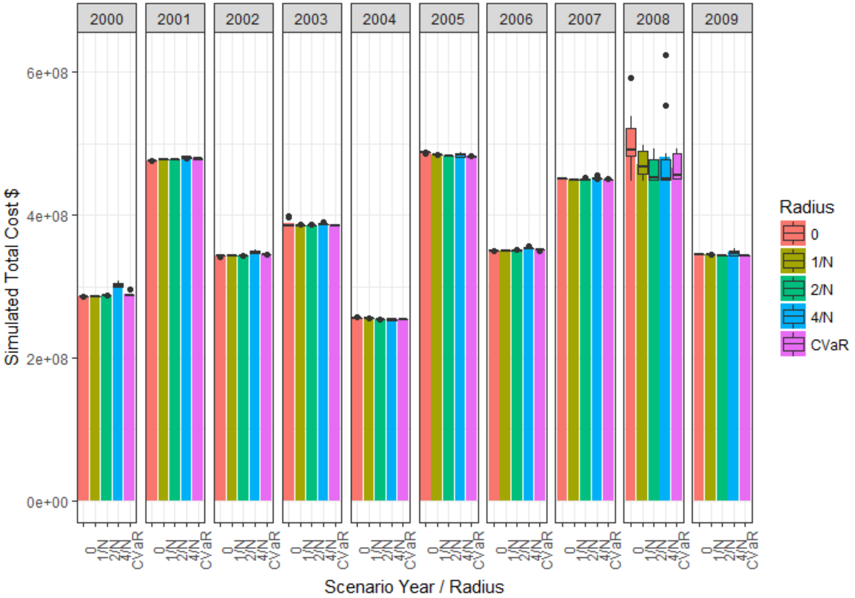

2000
2001
2002
2003
2004
2005
2006
2007
2008
2009
6e+08
4e+08
2e+08
0e+00
Simulated Total Cost $
0
1/N
2/N
4/N
CVaR
Scenario Year / Radius
Radius
0
1/N
2/N
4/N
CVaR

State Dependence

- At the least favourable state, $r = \bar{r}$.
- At the best state, $r = R$.
- In between, $\bar{r} \leq r(x) \leq R$, but $r(x)$ will be defined implicitly:

$$\rho(Z(x)) = \rho_{\bar{r}}(Z(x)) + \lambda(v^\top x)(\rho_{\bar{R}}(Z(x)) - \rho_{\bar{r}}(Z(x)))$$

where we assume that λ is a convex decreasing function of x.

Example I

Example

Suppose $z = \begin{bmatrix} 1 \\ 2 \\ 3 \end{bmatrix}$ (where the primal problem is a minimisation) and let $r = 1/3$.

Then

$$m = 3;\ \bar{z} = 2;\ s = std(z) = 0.816$$

and

$$p = \begin{bmatrix} 1/3 \\ 1/3 \\ 1/3 \end{bmatrix} + \left(\begin{bmatrix} 1 \\ 2 \\ 3 \end{bmatrix} - \begin{bmatrix} 2 \\ 2 \\ 2 \end{bmatrix} \right) \div \sqrt{3} \times \frac{1/3}{0.816} = \begin{bmatrix} 0.0976 \\ 0.333 \\ 0.569 \end{bmatrix}.$$

Example I

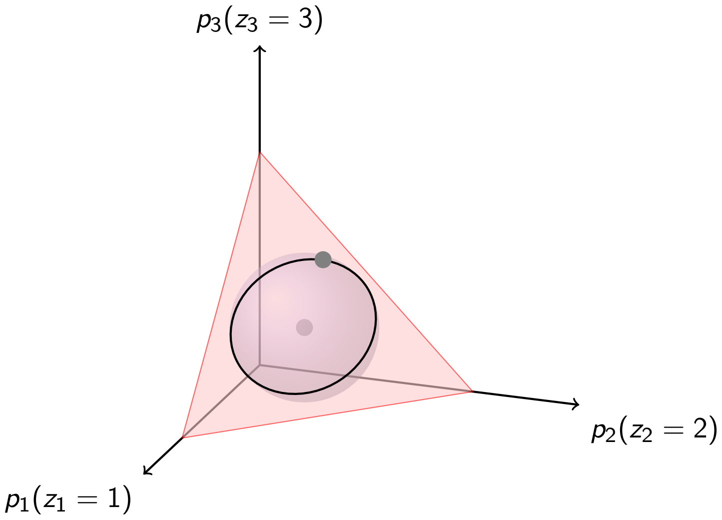

Example II

Example

Suppose $z = \begin{bmatrix} 1 \\ 2 \\ 5 \end{bmatrix}$ and let $r = 1/2 \geq \frac{1}{\sqrt{2 \times 3}}$.

Then

$$m = 3;\ \bar{z} = 8/3;\ s = std(z) = 1.7$$

and

$$p = \begin{bmatrix} 1/3 \\ 1/3 \\ 1/3 \end{bmatrix} + \left(\begin{bmatrix} 1 \\ 2 \\ 5 \end{bmatrix} - \begin{bmatrix} 8/3 \\ 8/3 \\ 8/3 \end{bmatrix} \right) \div \sqrt{3} \times \frac{1/2}{1.7} = \begin{bmatrix} 0.05026 \\ 0.2201 \\ 0.7296 \end{bmatrix}.$$

Larger Radii

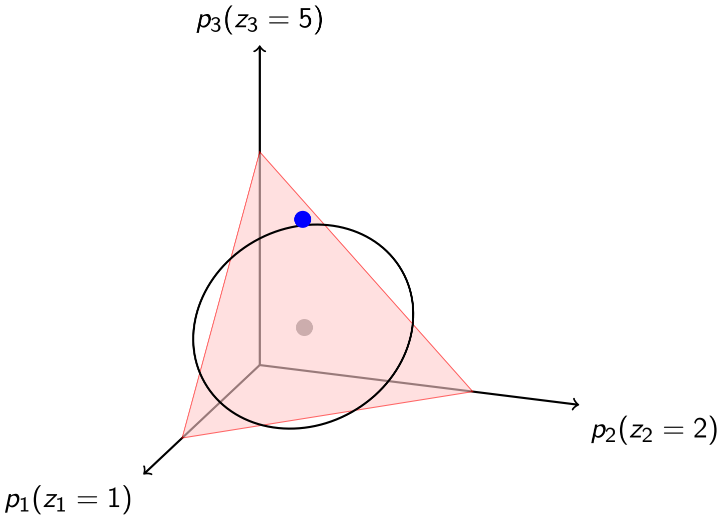

Example III

Example

Suppose $z = \begin{bmatrix} 1 \\ 2 \\ 3 \end{bmatrix}$ and let $r = 2/3$.

Then

$$m = 3; \;\; \bar{z} = 2; \;\; s = std(z) = 0.816$$

and

$$p = \begin{bmatrix} 1/3 \\ 1/3 \\ 1/3 \end{bmatrix} + \left(\begin{bmatrix} 1 \\ 2 \\ 3 \end{bmatrix} - \begin{bmatrix} 2 \\ 2 \\ 2 \end{bmatrix} \right) \div \sqrt{3} \times \frac{2/3}{0.816} = \begin{bmatrix} -0.138 \\ 0.333 \\ 0.805 \end{bmatrix}.$$

Example III

Example

Fix $p_1 = 0$ and reduce the problem to two dimensions. Now $z = \begin{bmatrix} 2 \\ 3 \end{bmatrix}$ and $m - k = 2$; $\bar{z} = 2.5$; $s = std(z) = 0.5$.

$$p_{2D} = \begin{bmatrix} 1/2 \\ 1/2 \end{bmatrix} + \left(\begin{bmatrix} 2 \\ 3 \end{bmatrix} - \begin{bmatrix} 2.5 \\ 2.5 \end{bmatrix} \right) \div 0.5 \times \sqrt{2\left(\frac{2}{3}\right)^2 - \frac{1}{3}} \div 2 = \begin{bmatrix} 0.127 \\ 0.873 \end{bmatrix}.$$

So

$$p = \begin{bmatrix} 0 \\ 0.127 \\ 0.873 \end{bmatrix}.$$

References I

- J. Bezanson, A. Edelman, S. Karpinski, and V.B. Shah, *Julia: A fresh approach to numerical computing*, SIAM Review **59** (2017), no. 1, 65–98.
- I. Dunning, J. Huchette, and M. Lubin, *JuMP: A modeling language for mathematical optimization*, SIAM Review **59** (2017), no. 2, 295–320.
- O. Dowson, *SDDP in Julia*, Tech. report, University of Auckland, 2017.
- M.V.F. Pereira and L.M.V.G. Pinto, *Multi-stage stochastic optimization applied to energy planning*, Mathematical Programming **52** (1991), no. 1-3, 359–375.