Discretization Issues

February 17, 2015

Position of the Problem...

We want to solve a closed-loop stochastic optimization problem, that is, a problem such that the decision variable $\boldsymbol{U}$ is a random variable which satisfies conditions imposed by the information structure of the problem defined by the random variable $\boldsymbol{Y}$.

We assume that the problem is dual effect free, even if we have to restrict the original admissible set $\mathcal{U}^{\mathrm{ad}} = \{\boldsymbol{U} \in \mathcal{U} \,,\ \ \boldsymbol{U} \preceq \boldsymbol{Y}\}$ of the problem to the "no dual effect" subset $\mathcal{U}^{\mathrm{nde}}$. Then, the information $\boldsymbol{Y}$ induced by the control $\boldsymbol{U}$ does not depend on $\boldsymbol{U}$.

We manipulate the measurability conditions from the algebraic point of view, that is, $\sigma(\boldsymbol{U}) \subset \sigma(\boldsymbol{Y})$.[14] In order to numerically solve the optimization problem, we have to approximate these constraints by using a finite representation.

[14] We turned back and use again the concept of σ-field rather than the one of π-field. Remember that these two notions coincide in the finite case.

and Problem under Consideration

The standard form of the problem we are interested in is

$$\mathcal{V}(\boldsymbol{W}, \mathcal{B}) = \min_{\boldsymbol{U} \in \mathcal{U}} \mathbb{E}\big(j(\boldsymbol{U}, \boldsymbol{W})\big) ,$$

subject to

$$\boldsymbol{U} \text{ is } \mathcal{B}\text{-measurable} ,$$

where $\mathcal{B} = \sigma(\boldsymbol{Y})$ is a fixed σ-field.

In order to obtain a numerically tractable approximation of this problem, we have to approximate

- the noise $\boldsymbol{W}$ by a "finite" noise $\boldsymbol{W}_n$ (Monte Carlo,...),
- the σ-field $\mathcal{B}$ by a "finite" σ-field $\mathcal{B}_n$ (partition,...).

Question: $\mathcal{V}(\boldsymbol{W}_n, \mathcal{B}_n) \longrightarrow \mathcal{V}(\boldsymbol{W}, \mathcal{B})$?

A Specific Instance of the Problem

A specific instance of the problem is the one which incorporates dynamical systems, that is, the stochastic optimal control problem:

$$\min_{(\boldsymbol{U}_0,\dots,\boldsymbol{U}_{T-1},\boldsymbol{X}_0,\dots,\boldsymbol{X}_T)} \mathbb{E}\left(\sum_{t=0}^{T-1} L_t(\boldsymbol{X}_t, \boldsymbol{U}_t, \boldsymbol{W}_{t+1}) + K(\boldsymbol{X}_T) \right)$$

subject to

$$\begin{aligned} \boldsymbol{X}_0 &= f_{-1}(\boldsymbol{W}_0) \ , \\ \boldsymbol{X}_{t+1} &= f_t(\boldsymbol{X}_t, \boldsymbol{U}_t, \boldsymbol{W}_{t+1}) \ , \quad t = 0, \dots, T-1 \ , \\ \boldsymbol{U}_t &\preceq \boldsymbol{Y}_t \ , \quad t = 0, \dots, T-1 \ . \end{aligned}$$

Assuming that $\sigma(\boldsymbol{Y}_t)$ are fixed σ-fields, a widely used approach to discretize this optimization problem is the so-called scenario tree method. We present it now, before considering the general case.

Lecture Outline

A Standard Stochastic Optimal Control Problem

Consider the following stochastic optimal control problem with a static (non-anticipative) information structure.

$$\min_{(\boldsymbol{U}_0,\dots,\boldsymbol{U}_{T-1},\boldsymbol{X}_0,\dots,\boldsymbol{X}_T)} \mathbb{E}\left(\sum_{t=0}^{T-1} L_t(\boldsymbol{X}_t, \boldsymbol{U}_t, \boldsymbol{W}_{t+1}) + K(\boldsymbol{X}_T) \right)$$

subject to

$$\boldsymbol{X}_0 = f_{-1}(\boldsymbol{W}_0) \;,$$
$$\boldsymbol{X}_{t+1} = f_t(\boldsymbol{X}_t, \boldsymbol{U}_t, \boldsymbol{W}_{t+1}) \;, \quad t = 0, \dots, T-1 \;,$$

$$\boldsymbol{U}_t \preceq h_t(\boldsymbol{W}_0, \dots, \boldsymbol{W}_t) \;, \qquad t = 0, \dots, T-1 \;.$$

Almost sure constraints (e.g. bound constraints on $\boldsymbol{X}_t$ and $\boldsymbol{U}_t$) may also be present in the formulation.

Scenario Tree Methodology

Obtain a finite dimensional approximation of the problem.

1. Discretize the noise process $\{\boldsymbol{W}_t\}$ using a scenario tree.
2. Copy out the measurability constraints on this structure:
   $$\boldsymbol{U}_t \preceq h_t(\boldsymbol{W}_0, \ldots, \boldsymbol{W}_t).$$
3. Write the dynamics and cost functions at the tree nodes:
   $$\boldsymbol{X}_{t+1} = f_t(\boldsymbol{X}_t, \boldsymbol{U}_t, \boldsymbol{W}_{t+1}).$$
4. Solve the problem using adequate mathematical programming techniques.

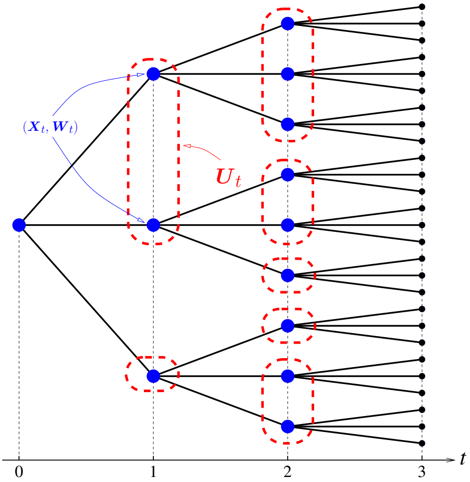

1. Discretize the Random Inputs

The tree architecture is characterized by the fact that each node of the tree corresponds to a unique past noise history but is generally followed by several possible future histories.

The tree is obtained by repeatedly using a finite approximation of the conditional probability laws $\mathbb{P}(\boldsymbol{W}_t \mid \boldsymbol{W}_0, \ldots, \boldsymbol{W}_{t-1})$:

$$\mathbb{P}(\boldsymbol{W}_0) \approx \{w_0^1, \ldots, w_0^{n_0}\} \;\rightsquigarrow\; \mathbb{P}(\boldsymbol{W}_1 \mid \boldsymbol{W}_0 = w_0^i) \approx \{w_1^{i,1}, \ldots, w_0^{i,n_1}\} \;\ldots$$

Note that this discretization scheme is much more sophisticated than the standard Monte Carlo sampling of $(\boldsymbol{W}_0, \ldots, \boldsymbol{W}_T)$.

The starting point may be a given collection of scenarios from which one constructs a tree by grouping the scenarios according to their (approximate) common past.

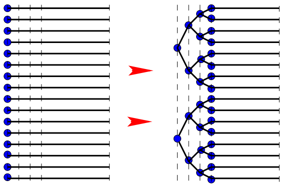

2. Copy out the Measurability Constraints

Assume that the information consists of the exact observation of all past noises: $\boldsymbol{Y}_t = (\boldsymbol{W}_0, \ldots, \boldsymbol{W}_t)$. Then, a different decision has to be attached at each node of the scenario tree.

But the method can face more general situations by grouping nodes of the scenario tree in order to represent the information structure induced by the $h_t(\boldsymbol{W}_0, \ldots, \boldsymbol{W}_t)$'s. For example, the information structure $h_t(\boldsymbol{W}_0, \ldots, \boldsymbol{W}_t) = \big(\hbar_0(\boldsymbol{W}_0), \ldots, \hbar_t(\boldsymbol{W}_t)\big)$ leads to a grouping of scenario tree nodes at each time step t, and ultimately produces a tree structure called the decision tree.

In all cases, the information structure is entirely coded within the scenario tree by means of those groups of nodes (one decision for each node of the decision tree in the previous example).

3. Write the Dynamics and Cost Functions

Consider a node $\nu \in \mathcal{N}$ of the scenario tree at time t, and denote:

- $\mathfrak{f}(\nu)$ the predecessor of node ν $(= \mu)$,
- $\pi(\nu)$ the probability of node ν,
- $\theta(\nu)$ the time index of node ν $(= t)$,
- $\gamma(\nu)$ the control index of node ν.

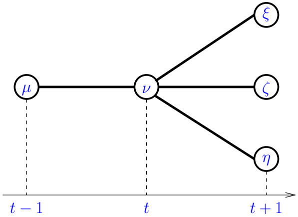

Note that the probability function π satisfies the following conditions:

$$\pi(\nu) = \sum_{\xi \in \mathfrak{f}^{-1}(\nu)} \pi(\xi) \quad , \quad \sum_{\nu \in \theta^{-1}(t)} \pi(\nu) = 1 \; .$$

Then, the dynamic equation from node μ to node ν writes

$$x_\nu = f_{\theta(\mathfrak{f}(\nu))}\big(x_{\mathfrak{f}(\nu)}, u_{\gamma(\mathfrak{f}(\nu))}, w_\nu\big) \; .$$

The cost induced by the transition is: $L_{\theta(\mathfrak{f}(\nu))}\big(x_{\mathfrak{f}(\nu)}, u_{\gamma(\mathfrak{f}(\nu))}, w_\nu\big)$.

4. Solve the Approximated Problem

The initial stochastic optimization problem boils down to

$$\min \left(\sum_{\nu \in \mathcal{N} \setminus \theta^{-1}(0)} \pi(\nu) L_{\theta(\mathfrak{f}(\nu))}\big(x_{\mathfrak{f}(\nu)}, u_{\gamma(\mathfrak{f}(\nu))}, w_\nu\big) + \sum_{\nu \in \theta^{-1}(T)} \pi(\nu) K(x_\nu) \right) ,$$

subject only to the dynamics constraints

$$x_\nu = f_{-1}(w_\nu) \qquad \forall \nu \in \theta^{-1}(0) ,$$

$$x_\nu = f_{\theta(\mathfrak{f}(\nu))}\big(x_{\mathfrak{f}(\nu)}, u_{\gamma(\mathfrak{f}(\nu))}, w_\nu\big) \qquad \forall \nu \in \mathcal{N} \setminus \theta^{-1}(0) .$$

The initial infinite dimensional stochastic optimization problem is approximated by a finite dimensional deterministic problem, that can be solved using relevant mathematical programming tools.

Note that this approximation corresponds to an optimal control problem with an arborescent (rather than linear) time structure.

Facts and Questions about the Scenario Tree Method

Dual Effect: it is mandatory that no dual effect holds true.

White noise: the noise process $(\boldsymbol{W}_0, \ldots, \boldsymbol{W}_T)$ may be correlated.

Perfect memory : this property is not required although useful.

Complexity: the amount of scenarios needed to achieve a given accuracy grows exponentially w.r.t. the number of time steps T of the problem (see [Shapiro, 2006]).

Tree structure: how to build a tree which is at the same time representative of the problem and numerically tractable?

Extrapolation: how to obtain feedback laws once the optimal decisions on the nodes of the scenario tree have been computed?

A huge literature is available on the scenario tree method. . .

Compact View of the Scenario Tree Approach

The previous stochastic optimal control problem depends on both a noise process $\boldsymbol{W}$ and a sequence of σ-fields $\mathcal{B}$. It can thus be represented under the compact form:

$$\mathcal{V}(\boldsymbol{W}, \mathcal{B}) = \min_{\boldsymbol{U} \in \mathcal{U}} \Big\{ \mathbb{E}\big(j(\boldsymbol{U}, \boldsymbol{W})\big) \quad \text{s.t.} \quad \boldsymbol{U} \;\; \mathcal{B}\text{-measurable} \Big\} ,$$

where $\mathcal{B} = \sigma\big(h(\boldsymbol{W})\big)$: SIS information structure.

The aim of the scenario tree method is to

- approximate the noise $\boldsymbol{W}$ by a "finite" noise $\boldsymbol{W}_n$,
- and deduce the approximated information $\mathcal{B}_n = h(\boldsymbol{W}_n)$.

In this framework, only one approximation is performed to obtain the approximated solution $\mathcal{V}\big(\boldsymbol{W}_n, h(\boldsymbol{W}_n)\big)$.

Such an approximation scheme converges to $\mathcal{V}(\boldsymbol{W}, \mathcal{B})$ (see [14]). But remember that the noise is discretized in a very specific way...

A simple SOC problem

$$\min_{\boldsymbol{U} \preceq \boldsymbol{W}_0} \mathbb{E}\left(\varepsilon \boldsymbol{U}^2 + (\boldsymbol{W}_0 + \boldsymbol{U} + \boldsymbol{W}_1)^2\right)$$

- The noises $(\boldsymbol{W}_0, \boldsymbol{W}_1)$ are independent random variables, each with a uniform probability distribution over $[-1, 1]$.
- The decision variable $\boldsymbol{U}$ is measurable w.r.t. $\boldsymbol{W}_0$: $\boldsymbol{U} \preceq \boldsymbol{W}_0$.
- The initial state is $\boldsymbol{X}_0 = \boldsymbol{W}_0$.
- The final state is $\boldsymbol{X}_1 = \boldsymbol{X}_0 + \boldsymbol{U} + \boldsymbol{W}_1$.

The goal is to minimize the expectation of $\left(\varepsilon \boldsymbol{U}^2 + \boldsymbol{X}_1^2\right)$, where ε is a "small" positive number (cheap control).

Note that this example exactly matches the Markovian setting.

Exact Solution of the Problem

$$\mathbb{E}\big(\varepsilon \boldsymbol{U}^2 + (\boldsymbol{W}_0 + \boldsymbol{U} + \boldsymbol{W}_1)^2\big) =$$
$$\mathbb{E}\big(\underbrace{\boldsymbol{W}_0^2}_{1/3} + \underbrace{\boldsymbol{W}_1^2}_{1/3} + (1+\varepsilon)\boldsymbol{U}^2 + 2\boldsymbol{U}\boldsymbol{W}_0 + 2\underbrace{\boldsymbol{U}\boldsymbol{W}_1}_{0} + 2\underbrace{\boldsymbol{W}_0\boldsymbol{W}_1}_{0}\big)$$

- The problem is thus equivalent to

$$\min_{\boldsymbol{U} \preceq \boldsymbol{W}_0} \frac{2}{3} + \mathbb{E}\big((1+\varepsilon)\boldsymbol{U}^2 + 2\boldsymbol{U}\boldsymbol{W}_0\big) ,$$

- By the first order optimality condition, the optimal solution is

$$\boldsymbol{U}^\sharp = -\frac{\boldsymbol{W}_0}{1+\varepsilon} .$$

- The associated optimal cost is readily calculated to be

$$J^\sharp = \frac{1}{3} \times \frac{1+2\varepsilon}{1+\varepsilon} = \frac{1}{3} + \mathrm{O}(\varepsilon) .$$

Noises Discretization "a la Monte Carlo"

We crudely sample the optimization problem.

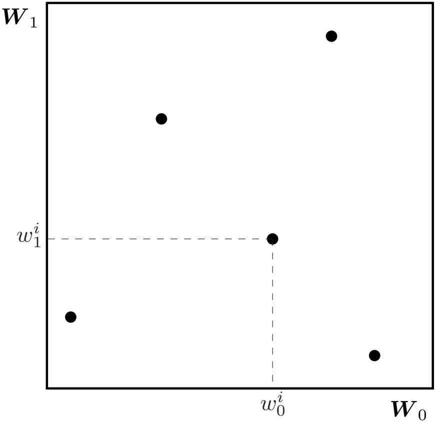

- To that purpose, we first consider a realization of a N-sample of the 2 noises $(\boldsymbol{W}_0, \boldsymbol{W}_1)$, that is, points in the square $\Omega = [-1, 1]^2$:

$$\left\{ (w_0^i, w_1^i) \right\}_{i=1,\ldots,N}$$

- The sample is first used to approximate the noise by the Monte Carlo method.

Discretized Information Structure

- We consider the N realizations $\{u^i\}_{i=1,\dots,N}$ of the decision variable $\boldsymbol{U}$, corresponding to the discretization of the noise,
- and we have to keep in mind that $\boldsymbol{U}$ should be measurable w.r.t. the first component $\boldsymbol{W}_0$ of the noise:

$$\boldsymbol{U} \preceq \boldsymbol{W}_0 \ .$$

- To translate this last condition in our discrete framework, we impose the constraint

$$\forall (i,j) \in \{1,\dots,N\}^2 \ , \ \ w_0^i = w_0^j \implies u^i = u^j \ ,$$

which prevents $\boldsymbol{U}$ from taking different values whenever the discretized representation of the noise $\boldsymbol{W}_0$ assumes the same value for two sample trajectories.

The Measurability Constraint is Not Effective!

The expression of the cost after discretization is

$$\frac{1}{N}\left(\sum_{i=1}^{N} \varepsilon\big(u^i\big)^2 + \big(w_0^i + u^i + w_1^i\big)^2\right) ,$$

and it is minimized w.r.t. $(u^1, \dots, u^N)$ under the constraints

$$u^i = u^j \quad \text{whenever} \quad w_0^i = w_0^j \; .$$

Since the N sample trajectories (w_0^i, w_1^i) of $(\boldsymbol{W}_0, \boldsymbol{W}_1)$ are produced by a Monte Carlo sampling over $[-1, 1]^2$, then, with probability 1,

$$w_0^i \neq w_0^j \qquad \forall (i, j) \quad \text{such that} \quad i \neq j \; .$$

Therefore, the above constraint is in fact never effective, so that the discretized expression of the cost is minimized independently for each individual sample i.

Something is Wrong. . .

The optimization problem associated to the i-th sample is

$$\min_{u_i \in \mathbb{R}} \varepsilon \big(u^i\big)^2 + \big(w_0^i + u^i + w_1^i\big)^2 \ ,$$

which yields the optimal value and the optimal cost

$$u_\flat^i = -\frac{w_0^i + w_1^i}{1+\varepsilon} \quad , \quad j_\flat^i = \varepsilon \frac{(w_0^i + w_1^i)^2}{1+\varepsilon} \ .$$

The averaged cost over the N samples is equal to

$$\frac{1}{N} \sum_{i=1}^{N} \frac{\varepsilon (w_0^i + w_1^i)^2}{1+\varepsilon} \underset{N \to +\infty}{\longrightarrow} \frac{2\,\varepsilon}{3\,(1+\varepsilon)} = \mathrm{O}(\varepsilon) \ .$$

This cost is far below the true optimal cost $J^\sharp = 1/3 + \mathrm{O}(\varepsilon)$!

However, any admissible solution (any $\boldsymbol{U}$ such that $\boldsymbol{U} \preceq \boldsymbol{W}_0$) cannot achieve a cost better than the optimal cost. . .

Real Value of the Discretized Problem Solution

- The resolution of the discretized problem derived from the previous Monte Carlo procedure yields optimal values

$$u_\flat^i = -\frac{w_0^i + w_1^i}{1+\varepsilon} ,$$

but not a random variable.

- The cost value of order ε is just a fake cost estimation, because we have not produced an admissible control.
- To evaluate the true cost of this "solution", we must first derive an admissible control for the initial problem, that is, a random variable $\boldsymbol{U}_\flat$ over $[-1, 1]^2$ with constant value along every vertical line of this square (since the horizontal axis represents the first component $\boldsymbol{W}_0$ of the noise, and since $\boldsymbol{U}_\flat$ has to be measurable with respect to $\boldsymbol{W}_0$ only).

Construction of an Admissible Control (1)

We assume that the sample points have been renumbered so that the first component w_0^i is increasing with i.

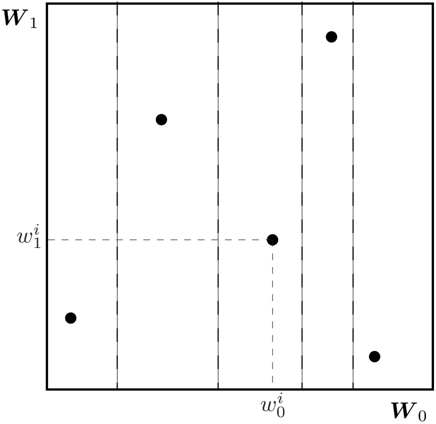

- Divide the square into N vertical strips by drawing vertical lines in the middle of segments $[w_0^i, w_0^{i+1}]$.
- The i-th strip is given by $[a^{i-1}, a^i] \times [-1, 1]$, with:

$$a^i = (w_0^i + w_0^{i+1})/2 \ ,$$

for $i = 2, \ldots, N-1$, ($a^0 = -1$ and $a^N = 1$).

Construction of an Admissible Control (2)

We define the solution $\boldsymbol{U}_\flat$ as the function of (w_0, w_1) which is piecewise constant over the square divided into those N strips, by using the optimal value $u_\flat^i$ in strip i:

$$\boldsymbol{U}_\flat(w_0, w_1) = \sum_{i=1}^{N} u_\flat^i \, \mathbf{1}_{[a^{i-1}, a^i] \times [-1,1]}(w_0, w_1) ,$$

where (w_0, w_1) ranges in the square $[-1, 1]^2$ and where $\mathbf{1}_A(\cdot)$ is the indicator function of the set A:

$$\mathbf{1}_A(x) = \begin{cases} 1 & \text{if } x \in A , \\ 0 & \text{otherwise} . \end{cases}$$

Of course, the control $\boldsymbol{U}_\flat$ depends on the N samples (w_0^i, w_0^i) by means of the values of the mid-points a^i's and the controls $u_\flat^i$'s.

Evaluation of the Expected Cost (1)

The corresponding cost value $\mathbb{E}\big(\varepsilon(\boldsymbol{U}_\flat)^2 + (\boldsymbol{W}_0 + \boldsymbol{U}_\flat + \boldsymbol{W}_1)^2\big)$ can be evaluated analytically (integration w.r.t. (w_0, w_1) over the square $[-1, 1]^2$), and is equal to

$$\frac{2}{3} + \sum_{i=1}^{N} \left((1+\varepsilon)\frac{a^i - a^{i-1}}{2} (u_\flat^i)^2 + \frac{(a^i)^2 - (a^{i-1})^2}{2} u_\flat^i \right) ,$$

where the values a^i and $u_\flat^i$ depend on the samples (w_0^i, w_1^i).

In order to assess the value of this estimate, we now compute its expectation when considering that the (w_0^i, w_1^i)'s are realizations of independent random variables $(\boldsymbol{W}_0^i, \boldsymbol{W}_1^i)$. This calculation is not straightforward because the w_0^i's have been reordered, so that we compute it numerically for different values of N.

Evaluation of the Expected Cost (2)

The cost provided by the admissible control $\boldsymbol{U}_\flat$ is estimated $2/3$.

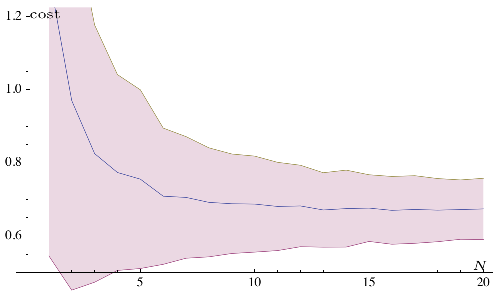

Figure: Estimated cost as a function of the number N of samples

This value neither corresponds to the true optimal cost ($1/3$), nor to the cost of the discrete problem (0). Moreover, the value $2/3$ is equal to that given by the best open-loop control: $\boldsymbol{U}_\star = 0$!

Scenario Tree Approach

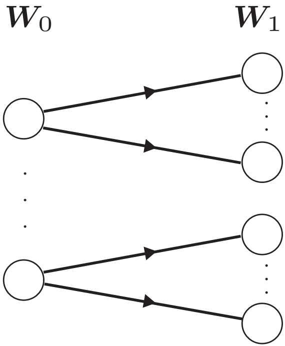

We consider $N_0 \times N_1$ scenarios

$$\left\{ \left(w_0^j, w_1^{jk} \right) \right\}_{j=1,\ldots,N_0}^{k=1,\ldots,N_1} .$$

- Notice that the discretization w_0^j of the first noise $\boldsymbol{W}_0$ only depends on $j = 1, \ldots, N_0$,
- whereas the discretization w_1^{jk} of the second noise $\boldsymbol{W}_1$ "hangs" from a given $j \in \{1, \ldots, N_0\}$ and then depends on $k = 1, \ldots, N_1$.

Scenario Tree Optimal Solution

On the scenario tree, the original cost $\mathbb{E}\big(\varepsilon \boldsymbol{U}^2 + (\boldsymbol{W}_0 + \boldsymbol{U} + \boldsymbol{W}_1)^2\big)$ is approximated by

$$\frac{1}{N_0}\sum_{j=1}^{N_0}\left(\varepsilon(u^j)^2 + \frac{1}{N_1}\sum_{k=1}^{N_1}(u^j + w_0^j + w_1^{jk})^2\right)$$

The solution of this approximated problem is

$$u_\natural^j = -\frac{w_0^j + \overline{w}_1^j}{1+\varepsilon} \ , \ \text{ where } \ \overline{w}_1^j = \frac{1}{N_1}\sum_{k=1}^{N_1} w_1^{jk} \ ,$$

to be compared with the naive Monte Carlo solution $u_\flat^j = -\dfrac{w_0^j + w_1^j}{1+\varepsilon}$.

Note that $\overline{w}_1^j$ can be interpreted as an estimate of the conditional expectation $\mathbb{E}\big(\boldsymbol{W}_1 \mid \boldsymbol{W}_0 = w_0^j\big)$.

Scenario Tree Optimal Cost

Let $(\overline{\sigma}_1^j)^2 = \frac{1}{N_1}\sum_{k=1}^{N_1}(w_1^{jk})^2$. The solution $u_\natural^j$ yields the cost

$$\frac{1}{N_0(1+\varepsilon)}\sum_{j=1}^{N_0}\left(\varepsilon(w_0^j)^2 + 2\varepsilon w_0^j \overline{w}_1^j - (\overline{w}_1^j)^2 + (1+\varepsilon)(\overline{\sigma}_1^j)^2\right) .$$

If we assume that the estimates $\overline{w}_1^j$ and $(\overline{\sigma}_1^j)^2$ converge towards their right values (respectively, 0 and $1/3$) as N_1 goes to infinity, then the scenario tree approach optimal cost gets close to

$$\frac{1}{N_0(1+\varepsilon)}\sum_{j=1}^{N_0}\left(\varepsilon(w_0^j)^2 + \frac{1+\varepsilon}{3}\right) \underset{N_0\to+\infty}{\longrightarrow} \frac{1}{3} + \mathrm{O}(\varepsilon) .$$

This cost is of the same order than the true optimal cost. However it does not correspond to an admissible solution. . .

Admissible Control and Associated Cost

As in the naive Monte Carlo method, we derive from the $u_\natural^j$'s an admissible solution $\boldsymbol{U}_\natural$ for the initial problem (piecewise constant fonction over N_0 strips of the square $[-1, 1]^2$). The cost provided by $\boldsymbol{U}_\natural$ is estimated $1/3$, corresponding to the true optimal cost.

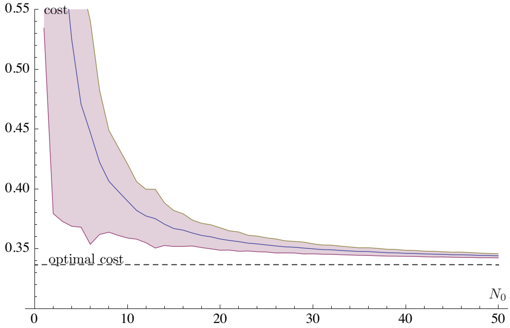

Figure: Estimated cost on a tree with N_0^2 scenarios

Where Do We Stand?

	True Solution	Naive Monte Carlo	Scenario Tree
Optimal Cost	$1/3 + \mathrm{O}(\varepsilon)$	$\mathrm{O}(\varepsilon)$	$1/3 + \mathrm{O}(\varepsilon)$
Feasible Control	$-\boldsymbol{W}_0/(1+\varepsilon)$	$-(w_0^i + w_1^i)/(1+\varepsilon)$	$-(w_0^i + \overline{w}_1^j)/(1+\varepsilon)$
Induced Cost	$1/3 + \mathrm{O}(\varepsilon)$	$2/3 + \mathrm{O}(\varepsilon)$	$1/3 + \mathrm{O}(\varepsilon)$

1. The naive Monte Carlo method
   - discretizes the noise process as a whole,
   - deduces the discretization of the measurability constraint,
   - yields a cost not better than the open-loop solution.
2. The scenario tree approach
   - discretizes the noise in a clever way (forward process),
   - deduces the discretization of the measurability constraint,
   - yields the optimal cost.

Hint: *the conditional probability laws are well estimated.*

Monte Carlo Interpretation of the Scenario Tree

$\boldsymbol{W}_0 \rightsquigarrow \{\boldsymbol{a}, \boldsymbol{b}, \boldsymbol{c}, \boldsymbol{d}\}$, $\boldsymbol{W}_1 \rightsquigarrow \{\{1\}, \{2, 3, 4\}, \{5, 6\}, \{7, 8, 9, 10\}\}$.

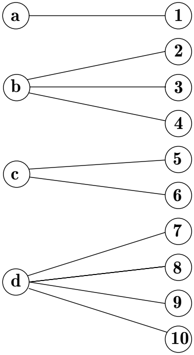

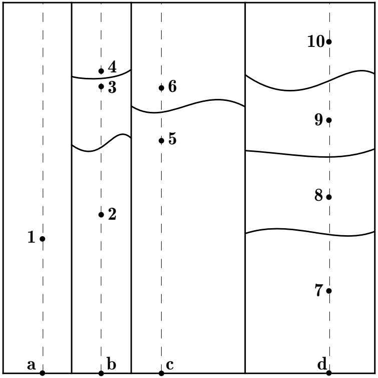

In a scenario tree, groups of samples are naturally aligned vertically!

Voronoi Quantization

However, others quantizations of Ω are possible.

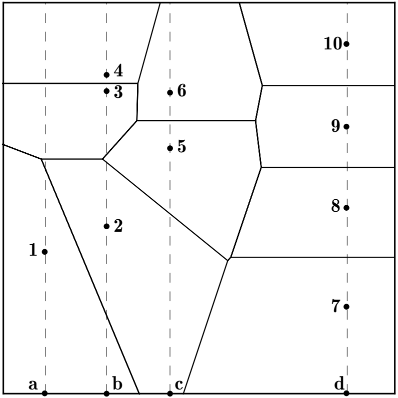

Given a set of points in the square $[-1, 1]^2$, the Voronoi tessellation minimizes the mean quadratic error among finite random variables taking given values. We in fact consider a discretized version of the random variable $(\boldsymbol{W}_0, \boldsymbol{W}_1)$, rather than a Monte Carlo sampling.

Independent Discretization of Noise and Information (1)

- Choose a discretization of the noise (8 cells).

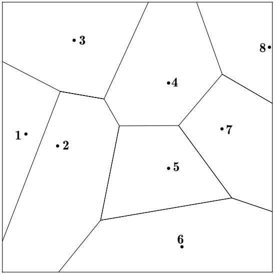

Noise

Independent Discretization of Noise and Information (2)

- Choose a discretization of the noise (8 cells).
- Choose a discretization of the information (5 cells).

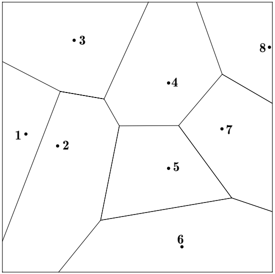

Noise

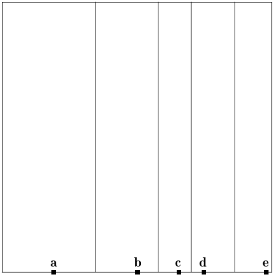

Information

Independent Discretization of Noise and Information (3)

- Choose a discretization of the noise (8 cells).
- Choose a discretization of the information (5 cells).
- Combine both discretizations (21 non empty cells).

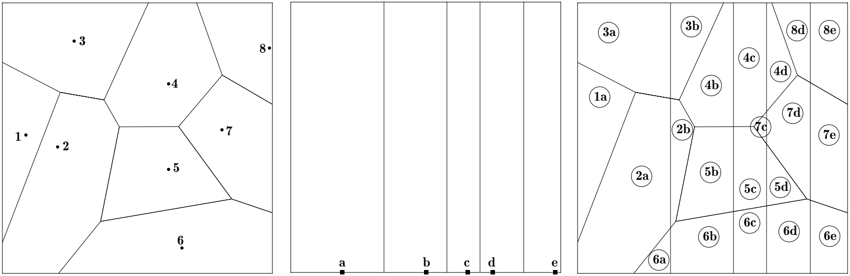

Noise Information Mixing

Independent Discretization of Noise and Information (4)

$\boldsymbol{W}_0 \rightsquigarrow \{a, b, c, d, e\}$, $\boldsymbol{W}_1 \rightsquigarrow \{1, 2, 3, 4, 5, 6, 7, 8\}$.

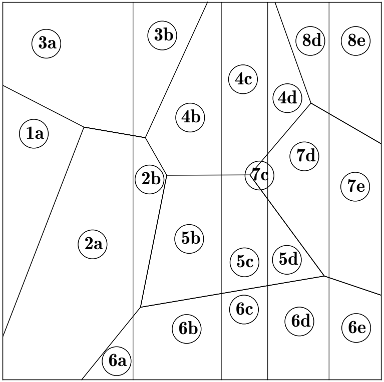

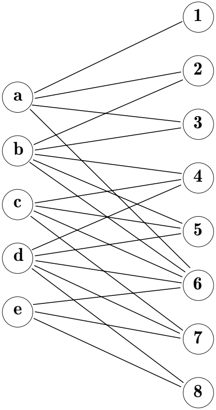

This approach does not necessarily produce a tree structure!

Discretized Optimization Problem

Using the notation $j(u, w_0, w_1) = \varepsilon u^2 + (w_0 + u + w_1)^2$, the discretized optimization problem is

$$\min_{\{u^k\}} \sum_{k \in \{\boldsymbol{a},\ldots,\boldsymbol{e}\}} \sum_{i=1}^{8} \pi^{ik} j(u^k, w_0^i, w_1^i) ,$$

where π^{ik} is the probability weight of the cell ik, u^k is the control value on the cell k and w^i the noise value on the cell i. Note that some of the π^{ik}'s are equal to zero.

The solution of this discretized problem can be computed (finite dimensional optimization). We expect that the optimal cost of the discretized problem converges to the true optimal cost $J^\sharp$ as the numbers of points in the 2 discrete sets associated to information and noise ($\{\boldsymbol{a}, \ldots, \boldsymbol{e}\}$ and $\{1, \ldots, 8\}$ in our example) go to infinity.

Problem and its Approximation

We consider the general form of a stochastic optimisation problem:

$$\mathcal{V}(\boldsymbol{W}, \mathcal{B}) = \min_{\boldsymbol{U} \in \mathcal{U}} \mathbb{E}\Big(j(\boldsymbol{U}, \boldsymbol{W})\Big) ,$$

subject to

$$\boldsymbol{U} \text{ is } \mathcal{B}\text{-measurable} .$$

We consider a sequence of random noises $\left\{\boldsymbol{W}_n\right\}_{n\in\mathbb{N}}$ and another sequence of σ-fields $\left\{\mathcal{B}_n\right\}_{n\in\mathbb{N}}$ such that the $\boldsymbol{W}_n$'s and the $\mathcal{B}_n$'s have "finite" representations, e.g.

- $\boldsymbol{W}_n = \sum_{i=1}^{n} w^i \mathbf{1}_{\Omega_i}$, $(\Omega_1, \ldots, \Omega_n)$ being a partition of Ω,
- $\mathcal{B}_n = \sigma(\Omega^1, \ldots, \Omega^n)$, $(\Omega^1, \ldots, \Omega^n)$ being a partition of Ω.

We are interested in the sequence of values $\left\{\mathcal{V}(\boldsymbol{W}_n, \mathcal{B}_n)\right\}_{n\in\mathbb{N}}$.

Convergence Notions for $\boldsymbol{W}$

These are rather standard notions.

- Convergence in distribution: $\boldsymbol{W}_n \xrightarrow{\mathcal{D}} \boldsymbol{W}$.

  $$\lim_{n\to+\infty} \mathbb{E}\Big(f(\boldsymbol{W}_n)\Big) = \mathbb{E}\Big(f(\boldsymbol{W})\Big) \text{ for all continuous bounded } f \ .$$

  This is the underlying concept in the Monte Carlo method: the empirical law defined by a N-sample $(\boldsymbol{W}^{(1)}, \ldots, \boldsymbol{W}^{(n)})$ of $\boldsymbol{W}$, that is, $\frac{1}{n}\sum_{i=1}^{n} \delta_{\boldsymbol{W}^{(i)}}$, weakly converges to $\mathbb{P}_{\boldsymbol{W}}$.

- Convergence in probability: $\boldsymbol{W}_n \xrightarrow{\mathbb{P}} \boldsymbol{W}$.

  $$\forall \epsilon > 0 \ , \quad \lim_{n\to+\infty} \mathbb{P}\Big(\big\|\boldsymbol{W}_n - \boldsymbol{W}\big\|_{\mathbb{W}} \geq \epsilon\Big) = 0 \ .$$

  This notion is much stronger than the previous one.

Convergence Notions for $\mathcal{B}$

These results are less known...

Strong Convergence of $\boldsymbol{\sigma}$-fields: $\mathcal{B}_n \to \mathcal{B}$.

$$\lim_{n \to +\infty} \mathbb{E}\Big(\boldsymbol{f} \;\Big|\; \mathcal{B}_n\Big) \xrightarrow{L^1} \mathbb{E}\Big(\boldsymbol{f} \;\Big|\; \mathcal{B}\Big) \text{ for all } \boldsymbol{f} \in \mathrm{L}^1(\mathbb{R}) .$$

Main properties.

1. The topology of the strong convergence is metrizable, so that the space $\mathcal{A}^\star$ of sub-fields of $\mathcal{A}$ is a complete separable metric space.
2. The σ-fields generated by a finite partition of Ω are dense in $\mathcal{A}^\star$ equipped with the previous metric.
3. Let $\{\boldsymbol{Y}_n\}_{n \in \mathbb{N}}$ be a sequence of random variables such that $\boldsymbol{Y}_n \xrightarrow{\mathbb{P}} \boldsymbol{Y}$ and $\sigma(\boldsymbol{Y}_n) \subset \sigma(\boldsymbol{Y}) \;\; \forall n$. Then, $\sigma(\boldsymbol{Y}_n) \to \sigma(\boldsymbol{Y})$.

Convergence Theorem

Theorem

Let $\mathcal{W} = L^q(\Omega, \mathcal{A}, \mathbb{P}; \mathbb{W})$ *and* $\mathcal{U} = L^r(\Omega, \mathcal{A}, \mathbb{P}; \mathbb{U})$, *with* $1 \leq q < +\infty$ *and* $1 \leq r < +\infty$. *Under the assumptions*

- $\boldsymbol{H_1}$ *the sequence* $\{\mathcal{B}_n\}_{n\in\mathbb{N}}$ *strongly converges to* $\mathcal{B}$, *and* $\mathcal{B}_n \subset \mathcal{B}$,
- $\boldsymbol{H_2}$ *the sequence* $\{\boldsymbol{W}_n\}_{n\in\mathbb{N}}$ *converges to* $\boldsymbol{W}$ *in* L^q*-norm,*
- $\boldsymbol{H_3}$ *the normal integrand* j *is such that*

$$\forall (u, u') \in \mathbb{U}^2 , \ \forall (w, w') \in \mathbb{W}^2 ,$$
$$\big| j(u, w) - j(u', w') \big| \leq \alpha \, \|u - u'\|_{\mathbb{U}}^r + \beta \, \|w - w'\|_{\mathbb{W}}^q ,$$

the convergence of the approximated optimal costs holds true

$$\lim_{n\to+\infty} \mathcal{V}\big(\boldsymbol{W}_n, \mathcal{B}_n\big) = \mathcal{V}\big(\boldsymbol{W}, \mathcal{B}\big) .$$

Using epi-convergence, it is possible to obtain the same results under weaker assumptions and to ensure the convergence of the sequence of the solutions.

Conclusions

- In the discretization of a SOC problem, there are two issues:
  - noise discretization,
  - information discretization.
- The naive Monte Carlo discretization provides a too weak convergence notion (in distribution, not in probability).
- The scenario tree methodology provides an effective way to discretize stochastic optimal control problem, but the two discretizations of information and of noise are bundled.
- Independent discretizations of noise and information offer
  - a greater latitude to select discretization schemes,
  - a way to obtain proper convergence results.

K. Barty, P. Carpentier, J.-P. Chancelier, G. Cohen, M. De Lara, and T. Guilbaud.
Dual effect free stochastic controls.
Annals of Operation Research, 142(1):41–62, February 2006.

H. S. Witsenhausen.
A counterexample in stochastic optimal control.
SIAM Journal on Control, 6(1):131–147, 1968.

Y. C. Ho and K. C. Chu.
Information structure in dynamic multi-person control problems.
Automatica, 10:341–351, 1974.

H. Kudo.
A note on the strong convergence of σ-algebras.
Annals of Probability, 2(1):76–83, 1974.

T. Pennanen.
Epi-convergent discretizations of multistage stochastic programs.
Mathematics of Operations Research, 30(1):245–256, 2005.

P. Carpentier, J.-P. Chancelier and M. De Lara.
Approximations of stochastic optimization problems subject to measurability constraints
SIAM Journal on Optimization, 19:1719–1734, 2009.

A. Shapiro.
On complexity of multistage stochastic programs.
Operations Research Letters, 34-1:1–8, 2006.

A. Shapiro, D. Dentcheva, and A. Ruszczyński.
Lectures on Stochastic Programming: Modeling and Theory.
SIAM, Philadelphia, 2009.