Introduction to Stochastic Optimization

Michel De Lara
Cermics, École des Ponts ParisTech
Université Paris-Est
France

École des Ponts ParisTech

November 16, 2014

Outline of the presentation

1. Working out classical examples
2. Framing stochastic optimization problems
3. Optimization with finite scenario space
4. Solving stochastic optimization problems by decomposition methods

Outline of the presentation

Working out classical examples

We will work out classical examples in Stochastic Optimization

- ▷ the blood-testing problem
  static, only risk
- ▷ the newsvendor problem
  static, only risk
- ▷ as a startup for stock management problems
  risk and time, with fixed information flow
- ▷ the secretary problem
  risk and time, with handleable information flow

Outline of the presentation

The blood-testing problem (R. Dorfman) is a static stochastic optimization problem

- A large number N of individuals are subject to a blood test
- The probability that the test is positive is p, the same for all people
- Individuals are stochastically independent
- The blood samples of k individuals are pooled and analyzed together
  - If the test is negative, this one test suffices for the k people
  - If the test is positive, each of the k persons must be tested separately, and $k+1$ tests are required, in all
- Find the value of k which minimizes the expected number of tests
- Find the minimal expected number of tests

In army practice, R. Dorfman achieved savings up to 80%

- For the first pool $\{1, \ldots, k\}$, the test is
  - negative with probability $(1-p)^k$ (by independence) $\to$ 1 test
  - positive with probability $1-(1-p)^k \to k+1$ tests
- When the pool size k is small compared to the number N of individuals, the blood samples $\{1, \ldots, N\}$ are split in approximately N/k groups, so that the expected number of tests is

$$J(k) \approx \frac{N}{k}[(1-p)^k + (k+1)(1-(1-p)^k)]$$

- For small p, the optimal solution is $k^\star \approx 1/\sqrt{p}$
- The minimal expected number of tests is about $J^\star \approx 2N\sqrt{p} < N$
- William Feller reports that, in army practice, R. Dorfman achieved savings up to 80%, compared to making N tests (take $p = 1/100$, giving $k^\star \approx 10$ and $J^\star \approx N/5$)

Outline of the presentation

The (single-period) newsvendor problem stands as a classic in stochastic optimization

- Traditionally known under the terminology "news*boy* problem", it is now coined the "news*vendor* problem" ;-)
- Each morning, the newsvendor must decide how many copies $u \in \mathbb{U} = \{0, 1, \ldots\}$ of the day's paper to order
- The newsvendor will meet an uncertain demand $w \in \mathbb{W} = \{0, 1, \ldots\}$
- The newsvendor faces an economic tradeoff
  - she pays the unitary purchasing cost c per copy, when she orders stock
  - she sells a copy at price p
  - if she remains with an unsold copy, it is worthless (perishable good)
- Therefore, the newsvendor's profit is uncertain,

$$\texttt{Payoff}(u, w) = -\underbrace{cu}_{\text{purchasing}} + \underbrace{p \min\{u, w\}}_{\text{selling}}$$

because it depends on the uncertain demand w

For you, Nature is rather random or hostile?

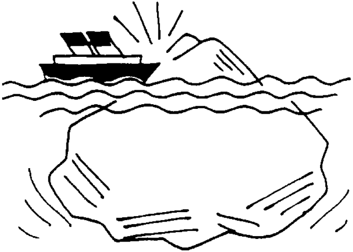

The newsvendor reveals her attitude towards risk in how she aggregates profit with respect to uncertainty

We formulate a problem of profit maximization

- In the robust or pessimistic approach,
  the newsvendor maximizes the worst payoff

$$\max_{u \in \mathbb{U}} \underbrace{\min_{w \in \overline{\mathbb{W}}} \texttt{Payoff}(u,w)}_{\text{worst payoff}}$$

  as if Nature were malevolent

- In the stochastic or expected approach, the newsvendor solves

$$\max_{u \in \mathbb{U}} \underbrace{\mathbb{E}_w[\texttt{Payoff}(u,w)]}_{\text{expected payoff}}$$

  as if Nature played stochastically

If the newsvendor maximizes the worse profit

- We suppose that
  - the demand w belongs to a set $\overline{\mathbb{W}} = [\![w^\flat, w^\sharp]\!]$
  - the newsvendor knows the set $[\![w^\flat, w^\sharp]\!]$
- The worse profit is

$$J(u) = \min_{w \in [\![w^\flat, w^\sharp]\!]} [-cu + p\min\{u, w\}] = -cu + p\min\{u, w^\flat\}$$

- Show that the order $u^\star = w^\flat$ maximizes the above expression $J(u)$
- Once the newsvendor makes the optimal order $u^\star = w^\flat$,
  the optimal profit is $w \mapsto (p - c)w^\flat$
  which, here, is no longer uncertain

If the newsvendor maximizes the expected profit

- We suppose that
  - the demand w is a random variable
  - the newsvendor knows the probability distribution $\mathbb{P}$ of w

$$\pi_0 = \mathbb{P}(w = 0), \ \pi_1 = \mathbb{P}(w = 1) \ \ldots$$

- The expected profit is

$$J(u) = \mathbb{E}_w[-cu + p \min\{u, w\}] = -cu + p\mathbb{E}[\min\{u, w\}]$$

- Find an order $u^\star$ which maximizes the above expression $J(u)$
  - by calculating $J(u + 1) - J(u)$
  - then using the *decumulative distribution function* $d \mapsto \mathbb{P}(w > d)$

Here stand some steps of the computation

$$
\begin{aligned}
J(u) &= -cu + p\mathbb{E}[\min\{u, w\}] \\
\min\{u, w\} &= u\mathbf{1}_{u<w} + w\mathbf{1}_{u\geq w} \\
\min\{u+1, w\} &= (u+1)\mathbf{1}_{u+1\leq w} + w\mathbf{1}_{u+1>w} \\
&= (u+1)\mathbf{1}_{u<w} + w\mathbf{1}_{u\geq w} \\
\min\{u+1, w\} - \min\{u, w\} &= \mathbf{1}_{u<w}
\end{aligned}
$$

$$J(u+1) - J(u) = -c + p\mathbb{E}[\mathbf{1}_{u<w}] = -c + p\mathbb{P}(w > u) \downarrow \text{ with } u$$

- An optimal decision $u^\star$ satisfies

$$\mathbb{P}(w > u^\star) \approx \frac{c}{p} = \frac{\text{cost}}{\text{price}}$$

- Once the newsvendor makes the optimal order $u^\star$, the optimal profit is the random variable $w \mapsto -cu^\star + p\min\{u^\star, w\}$

Where do we stand after having worked out two examples?

- When you move from deterministic optimization to optimization under uncertainty, you come accross the issue of risk attitudes
- Risk attitudes materialize in the a priori knowledge on the uncertainties
  - either probabilistic/stochastic
    - independence and Bernoulli distributions in the blood test example
    - uncertain demand faced by the newsvendor modeled as a random variable
  - or set-membership
    - uncertain demand faced by the newsvendor modeled by a set

Where do we stand after having worked out two examples?

- When you move from deterministic optimization to optimization under uncertainty, you come accross the issue of risk attitudes
- Risk attitudes materialize in the a priori knowledge
  on the uncertainties
  - either probabilistic/stochastic
    - independence and Bernoulli distributions in the blood test example
    - uncertain demand faced by the newsvendor modeled as a random variable
  - or set-membership
    - uncertain demand faced by the newsvendor modeled by a set
- In addition, when you make a succession of decisions, you need to specify what you know (of the uncertainties) before each decision, and what you know before each decision may depend or not on your previous actions
- Let us turn to the inventory problem

Outline of the presentation

Inventory control dynamical model

Consider the control dynamical model

$$x(t+1) = x(t) + u(t) - w(t)$$

where

- time $t \in \{t_0, \ldots, T\}$ is discrete (days, weeks or months, etc.)
- $x(t)$ is the stock at the beginning of period t, belonging to $\mathbb{X} = \mathbb{Z} = \{\ldots, -2, -1, 0, 1, 2, \ldots\}$
- $u(t)$ is the stock ordered at the beginning of period t, belonging to $\mathbb{U} = \mathbb{N} = \{0, 1, 2, \ldots\}$
- $w(t)$ is the uncertain demand during the period t, belonging to $\mathbb{W} = \mathbb{N}$

(When $x(t) < 0$, this corresponds to a *backlogged demand*,
supposed to be filled immediately once inventory is again available)

Inventory optimization criterion

- The costs incurred in period t are
  - purchasing costs: $cu(t)$
  - shortage costs: $b\max\{0, -\big(x(t)+u(t)-w(t)\big)\}$
  - holding costs: $h\max\{0, x(t)+u(t)-w(t)\}$
- On the period from t_0 to T, the costs sum up to

$$\sum_{t=t_0}^{T-1}\Big[\underbrace{cu(t)}_{\text{purchasing}} + \overbrace{\underbrace{b\max\{0,-\big(x(t)+u(t)-w(t)\big)\}}_{\text{shortage}} + \underbrace{h\max\{0,x(t)+u(t)-w(t)\}}_{\text{holding}}}^{\texttt{Cost}(x(t)+u(t)-w(t))}\Big]$$

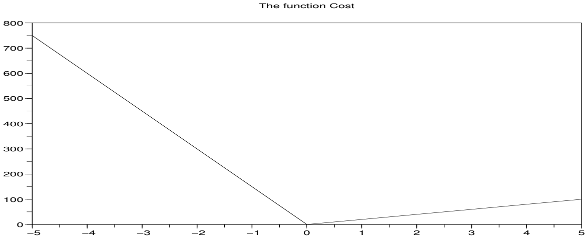

Probabilistic assumptions and risk neutral formulation of the inventory stochastic optimization problem

- We suppose that the sequence of demands $w(t_0), \ldots, w(T-1)$ is a stochastic process with distribution $\mathbb{P}$
- We consider the inventory sochastic optimization problem

$$\min_{u(\cdot)} \mathbb{E} \sum_{t=t_0}^{T-1} \left[cu(t) + \texttt{Cost}\big(x(t) + u(t) - w(t)\big) \right]$$

Information flow and closed-loop formulation of the inventory stochastic optimization problem

- Let $u(\cdot) = u(t_0), \ldots, u(T-1)$ and consider

$$\underbrace{\min_{u(\cdot)}}_{\text{meaning what?}} \quad \mathbb{E} \sum_{t=t_0}^{T-1} \big[c u(t) + \mathtt{Cost}\big(x(t) + u(t) - w(t)\big) \big]$$

- The decision $u(t)$ at time t belongs to the control set $\mathbb{U}$
- $u(t)$ is a random variable, like are all demands $w(t_0), \ldots, w(T-1)$
- and like are all states $x(t)$ by the dynamics $x(t+1) = x(t) + u(t) - w(t)$

We express that the decision $u(t)$ at time t depends on the past $w(t_0), \ldots, w(t)$

$$u(t) \text{ is measurable w.r.t. } \underbrace{\big(w(t_0), \ldots, w(t)\big)}_{\text{past}}$$

Where do we stand?

- In addition to risk, we have to pay attention to the information flow
- When we make a succession of decisions, we need to specify what we know (of the uncertainties) before each decision, and this information may depend or not on our previous actions
- Let us now turn to the secretary problem

Outline of the presentation

The secretary problem stands as a classic optimal stopping problem

- A firm has opened a single secretarial position to fill (or a princess will only accept one "fiancé")
- Secretary applicants (Alice, Bob, Claire, etc.) can be compared by their absolute rank, corresponding to his/her quality for the position (Alice is 7, Bob is 15, Claire has top rank 1, etc.)
- The interviewer does not know the absolute rank
- The interviewer screens N applicants one-by-one in random order (Bob, then Claire, then Alice, etc.)
- The interviewer is able to rank the applicants interviewed so far (for the job, Claire is better than Alice, who is better than Bob, etc.)
- After each interview, the interviewer decides
  - either to select the applicant (and the process stops)
  - or to reject the applicant (and the process goes on), knowing that, once rejected, an applicant cannot be recalled

Here, a strategy is a stopping rule

- There are N applicants for the position
- The value of N is known
- A strategy provides the number $\nu \in \{1, \ldots, N\}$ of applicants interviewed, as a fonction of the relative ranking of the applicants interviewed so far
- A stopping time is a random variable ν, such that, for any $n = 1, \ldots, N$, the event $\{\nu = n\}$ depends at most upon what happened before interview n
- The interviewer maximizes the probability to select the best applicant, among all strategies

Open-loop strategies yield a probability $1/N$

- An open-loop strategy does not use the information collected up to applicant n, except for the clock n
- Therefore, for any $n = 1, \ldots, N$, the event $\{\nu = n\}$ depends only on n, and not on what happened before interview n
- Thus, an open-loop strategy is a deterministic stopping time ν
- For instance, $\nu = 1$ (constant stopping time) is an open-loop strategy: you select the first applicant
- If you adopt the strategy $\nu = 1$, the probability of selecting the best applicant is $1/N$
- For a fixed $k \in \{1, \ldots, N\}$, the strategy $\nu = k$ also yields probability $1/N$

The best closed loop strategy yields a probability $\approx 1/e$

- A candidate is an applicant who, when interviewed, is better than all the applicants interviewed previously
- For a fixed $k \in \{1, \ldots, N\}$, consider the strategy ν_k:
  - select the first candidate popping up after k applicants have been interviewed
  - or select the last applicant N in case no candidate appears
- We will now show that, when the number N of applicants is large, the best among the strategies ν_k, $k = 1, \ldots, N$, is achieved for

$$k^\star \approx \frac{N}{e} \, , \quad \text{the so-called 37\% rule}$$

- The probability of selecting the best applicant is $\approx 1/e$

$$\underbrace{\frac{1}{e}}_{\text{closed loop}} > \underbrace{\frac{1}{N}}_{\text{open loop}}$$

Here stand some steps of the computation (1)

We denote $p(k)$ the probability to select the best applicant with strategy ν_k

$$
\begin{aligned}
p(k) &= \sum_{m=k}^{n} \mathbb{P}(\text{ applicant m is selected } \mid \text{ applicant m is the best }) \\
&\qquad \times \mathbb{P}(\text{ applicant m is the best }) \\
&= \sum_{m=k}^{n} \mathbb{P}(\text{ applicant m is selected } \mid \text{ applicant m is the best }) \times \frac{1}{n}
\end{aligned}
$$

- If applicant m is the best applicant, then m is selected if and only if the best applicant among the first $m-1$ applicants is among the first $k-1$ applicants that were rejected
- Deduce that, when $m \geq k$,

$$
\mathbb{P}(\text{ applicant m is selected } \mid \text{ applicant m is the best }) = \frac{k-1}{m-1}
$$

Here stand some steps of the computation (2)

- Sum over $m \geq k$ and obtain

$$p(k) = \sum_{m=k}^{n} \frac{k-1}{m-1} \times \frac{1}{n} = \frac{k-1}{n} \sum_{m=k}^{n} \frac{1}{m-1}$$

- Compute the difference

$$\begin{aligned} n[p(k+1) - p(k)] &= \sum_{m=k+1}^{n} \frac{1}{m-1} - 1 \\ &= \sum_{m=k+1}^{n} \frac{1}{m-1} - 1 \\ &\approx \log n - \log k - 1 \\ &= \log(\frac{n}{ke}) \end{aligned}$$

The optimal strategy is called the 37% rule

- ▷ What is the $k^\star$ that maximizes $p(k)$? The 37% rule:

$$k^\star \approx \frac{N}{e} \text{ where } \log e = 1$$

- ▷ What is $p(k^\star)$ when N runs to $+\infty$?

$$p(k^\star) \approx \frac{1}{e} \approx 37\%$$

Where do we stand after having worked out the secretary problem?

- In a stopping time problem, as long as you do not stop, you collect information
- This information is valuable for forthcoming decisions
- For Markov decision problems, information is condensed in a state
- Stochastic control problems display trade-off between exploration and exploitation

Many decision problems illustrate the trade-off between exploration and exploitation

- ▷ deciding where to dig
- ▷ animal foraging
- ▷ job search
- ▷ devoting resources to research

The interplay between information and decision makes stochastic control problems especially tricky and difficult

- Decision $\rightarrow$ information $\rightarrow$ decision $\rightarrow$ information $\rightarrow \cdots$
- Decisions generally induce a dual effect,
  a terminology which tries to convey the idea that present decisions have two, often conflicting, effects or objectives:
  - directly contributing to optimizing the cost function,
    on the one hand
  - modifying the future information available for forthcoming decisions,
    on the other hand
- Problems with dual effect are among the most difficult decision-making problems

Summary

- Stochastic optimization = risk + information
- Risk is in the eyes of the beholder ;-)
- Information can be either revealed progressively
  - in a fixed way
  - or depending on past decisions
- Now, we turn to the mathematical framing of stochastic optimization problems

Outline of the presentation

Let us work out a toy example of economic dispatch as a cost-minimization problem under supply-demand balance

- Production: consider two energy production units
  - a "cheap" limited one with which we can produce quantity q_0, with $0 \leq q_0 \leq q_0^\sharp$, at cost $c_0 q_0$
  - an "expensive" unlimited one with which we can produce quantity q_1, with $0 \leq q_1$, at cost $c_1 q_1$, with $c_1 > c_0$
- Consumption: the demand is $D \geq 0$
- Balance: ensuring at least the demand

$$D \leq q_0 + q_1$$

- Optimization: total costs minimization

$$\min_{q_0, q_1} \underbrace{c_0 q_0 + c_1 q_1}_{\text{total costs}}$$

When the demand D is deterministic, the optimization problem is well posed

- The deterministic demand D is a single number, and we consider

$$\min_{q_0,q_1} c_0 q_0 + c_1 q_1$$

$$\text{under the constraints} \quad \begin{array}{rl} 0 & \leq q_0 \leq q_0^{\sharp} \\ 0 & \leq q_1 \\ D & \leq q_0 + q_1 \end{array}$$

- The solution is $q_0^{\star} = \min\{q_0^{\sharp}, D\}\,, \quad q_1^{\star} = [D - q_0^{\sharp}]_{+}$, that is,
  - if the demand D is below the capacity $q_0^{\sharp}$ of the "cheap" energy source

$$D \leq q_0^{\sharp} \Rightarrow q_0^{\star} = D\,, \quad q_1^{\star} = 0$$

  - if the demand D is above the capacity $q_0^{\sharp}$ of the "cheap" energy source,

$$D > q_0^{\sharp} \Rightarrow q_0^{\star} = q_0^{\sharp}\,, \quad q_1^{\star} = D - q_0^{\sharp}$$

- Now, what happens when the demand D is no longer deterministic?

If we know the demand beforehand, the optimization problem is deterministic

- ▷ We suppose that the demand is a random variable $D : \Omega \to \mathbb{R}_+$
- ▷ If we solve the problem for each possible value $D(\omega)$ of the random variable D, when $\omega \in \Omega$, we obtain

$$q_0(\omega) = \min\{q_0^\sharp, D(\omega)\}\,, \quad q_1(\omega) = [D(\omega) - q_0^\sharp]_+$$

  and we face an informational issue
- ▷ Indeed, we treat the demand D as if it were observed before making the decisions q_0 and q_1
- ▷ When the demand D is not observed, how can we do?

What happens if we replace the uncertain value D of the demand by its mean $\overline{D}$ in the deterministic solution?

- If we suppose that the demand D is a random variable $D : \Omega \to \mathbb{R}_+$, with mathematical expectation $\mathbb{E}(D) = \overline{D}$
- and that we propose the "deterministic solution"

$$q_0^{(\overline{D})} = \min\{q_0^\sharp, \overline{D}\} \ , \ \ q_1^{(\overline{D})} = [\overline{D} - q_0^\sharp]_+$$

- we cannot assure the inequality

$$\underbrace{D(\omega)}_{\text{uncertain}} \leq \underbrace{q_0 + q_1}_{\text{deterministic}} \ , \ \ \forall \omega \in \Omega$$

because $\max_{\omega \in \Omega} D(\omega) > \overline{D} = q_0^{(\overline{D})} + q_1^{(\overline{D})}$

- Are there better solutions among the deterministic ones?

When the demand D is bounded above, the robust optimization problem has a solution

- In the robust optimization problem, we minimize

$$\min_{q_0, q_1} c_0 q_0 + c_1 q_1$$

$$\text{under the constraints} \quad \begin{aligned} 0 &\leq q_0 \leq q_0^{\sharp} \\ 0 &\leq q_1 \\ D(\omega) &\leq q_0 + q_1 \ , \quad \forall \omega \in \Omega \end{aligned}$$

- When $D^{\sharp} = \max_{\omega \in \Omega} D(\omega) < +\infty$, the solution is $q_0^{\star} = \min\{q_0^{\sharp}, D^{\sharp}\} \ , \quad q_1^{\star} = [D^{\sharp} - q_0^{\sharp}]_+$
- Now, the total cost $c_0 q_0^{\star} + c_1 q_1^{\star}$ is an increasing function of the upper bound $D^{\sharp}$ of the demand
- Is it not too costly to optimize under the worst-case situation?

Where do we stand?

- ▷ When the demand D is deterministic, the optimization problem is well posed
- ▷ If we know the demand beforehand, the optimization problem is deterministic
- ▷ If we replace the uncertain value D of the demand by its mean $\overline{D}$ in the deterministic solution, we remain with a feasability issue
- ▷ When the demand D is bounded above, the robust optimization problem has a solution, but it is costly

Where do we stand?

- When the demand D is deterministic, the optimization problem is well posed
- If we know the demand beforehand, the optimization problem is deterministic
- If we replace the uncertain value D of the demand by its mean $\overline{D}$ in the deterministic solution, we remain with a feasability issue
- When the demand D is bounded above, the robust optimization problem has a solution, but it is costly

To overcome the above difficulties, we propose to introduce stages

$$\underbrace{D(\omega)}_{\text{uncertain}} \leq \underbrace{q_0}_{\text{deterministic}} + \underbrace{q_1(\omega)}_{\text{uncertain}} , \quad \forall \omega \in \Omega$$

- the decision q_0 is made before observing the demand $D(\omega)$
- the decision $q_1(\omega)$ is made after observing the demand $D(\omega)$

To overcome the above difficulties, we turn to stochastic optimization

- We suppose that the demand D is a random variable, and minimize

$$\min_{q_0, q_1} \mathbb{E}[c_0 q_0 + c_1 q_1]$$

$$\text{under the constraints} \quad \begin{array}{ll} 0 & \leq q_0 \leq q_0^\sharp \\ 0 & \leq q_1 \\ D & \leq q_0 + q_1 \\ q_1 & \text{depends upon } D \end{array}$$

and we emphasize two issues, new with respect to the deterministic case

- expliciting online information issue:
  the decision q_1 depends upon the random variable D
- expliciting risk attitudes:
  we aggregate the total costs with respect to all possible values by taking the expectation $\mathbb{E}[c_0 q_0 + c_1 q_1]$

Turning to stochastic optimization forces one to specify online information

- We suppose that the demand D is a random variable, and minimize

$$\min_{q_0, q_1} \mathbb{E}[c_0 q_0 + c_1 q_1]$$

$$\text{under the constraints} \quad \begin{array}{ll} 0 & \leq q_0 \leq q_0^\sharp \\ 0 & \leq q_1 \\ D & \leq q_0 + q_1 \\ q_1 & \text{depends upon } D \end{array}$$

- specifying that the decision q_1 depends upon the random variable D, whereas q_0 does not, forces to consider two stages and a so-called non-anticipativity constraint (more on that later)
  - first stage: q_0 does not depend upon the random variable D
  - second stage: q_1 depends upon the random variable D

Turning to stochastic optimization forces one to specify risk attitudes

- We suppose that the demand D is a random variable, and minimize

$$\min_{q_0, q_1} \mathbb{E}[c_0 q_0 + c_1 q_1]$$

$$\text{under the constraints} \quad \begin{array}{ll} 0 & \leq q_0 \leq q_0^\sharp \\ 0 & \leq q_1 \\ D & \leq q_0 + q_1 \\ q_1 & \text{depends upon } D \end{array}$$

- Now that q_1 depends upon the random variable D,
  it is also a random variable, and so is the total cost $c_0 q_0 + c_1 q_1$;
  therefore, we have to aggregate the total costs with respect to all possible values, and we chose to do it by taking the expectation $\mathbb{E}[c_0 q_0 + c_1 q_1]$

In the uncertain framework, two additional questions must be answered with respect to the deterministic case

Question (expliciting risk attitudes)

How are the uncertainties taken into account
in the payoff criterion and in the constraints?

Question (expliciting available online information)

Upon which online information are decisions made?

Outline of the presentation

Water inflows historical scenarios

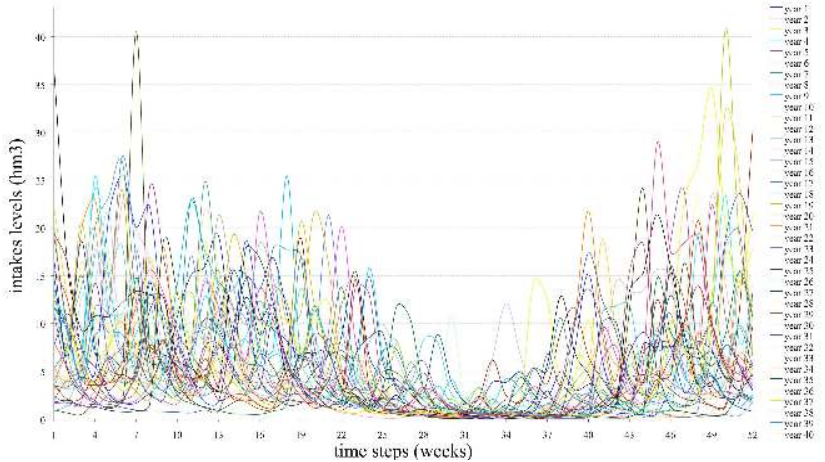

We call scenario a temporal sequence of uncertainties

Scenarios are special cases of "states of Nature"

A scenario (pathway, chronicle) is a sequence of uncertainties

$$w(\cdot) := \big(w(t_0), \dots, w(T-1)\big) \in \Omega := \mathbb{W}^{T-t_0}$$

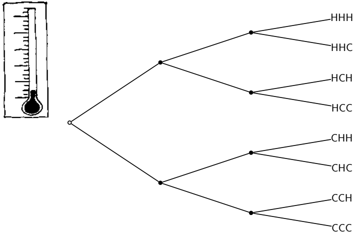

El tiempo se bifurca perpetuamente hacia innumerables futuros
(Jorge Luis Borges, *El jardín de senderos que se bifurcan*)

Beware! Scenario holds a different meaning in other scientific communities

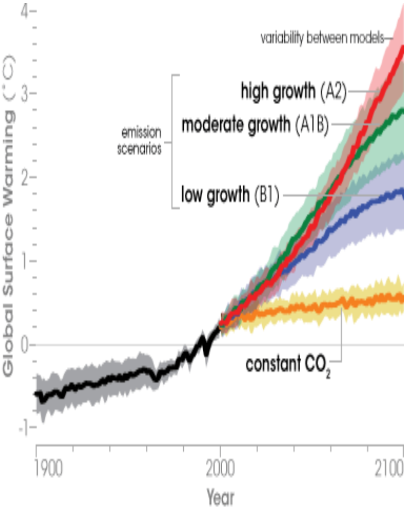

- In practice, what modelers call a "scenario" is a mixture of
  - a sequence of uncertain variables (also called a pathway, a chronicle)
  - a policy `Pol`
  - and even a static or dynamical model
- In what follows

  scenario = pathway = chronicle

Outline of the presentation

The output of a stochastic optimization problem is a random variable. How can we rank random variables?

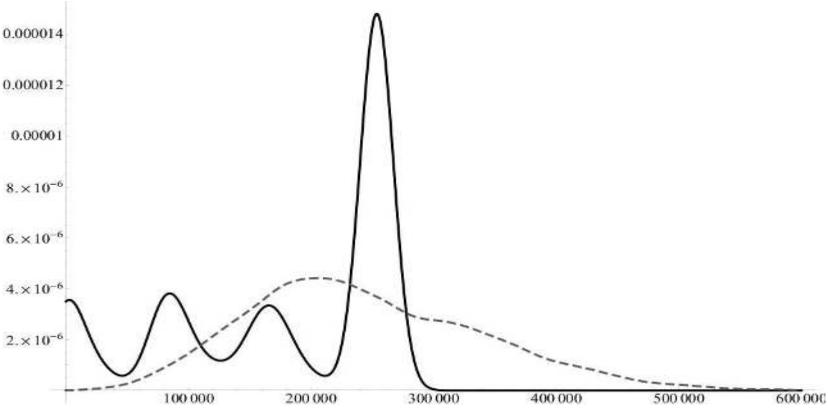

How are the uncertainties taken into account in the payoff criterion and in the constraints?

In a probabilistic setting, where uncertainties are random variables, a classical answer is

- to take the mathematical expectation of the payoff (risk-neutral approach)

$$\mathbb{E}(\text{payoff})$$

- and to satisfy all (physical) constrainsts almost surely that is, practically, for all possible issues of the uncertainties (robust approach)

$$\mathbb{P}(\text{constrainsts}) = 1$$

But there are many other ways to handle risk: robust, worst case, risk measures, in probability, almost surely, by penalization, etc.

A policy and a criterion yield a real-valued payoff

Given an admissible policy $\texttt{Pol} \in \mathcal{U}^{ad}$ and a scenario $w(\cdot) \in \Omega$, we obtain a payoff

$$\texttt{Payoff}\big(\texttt{Pol}, w(\cdot)\big)$$

Policies/Scenarios	$w^A(\cdot) \in \Omega$	$w^B(\cdot) \in \Omega$	...
$\texttt{Pol}_1 \in \mathcal{U}^{ad}$	$\texttt{Payoff}\big(\texttt{Pol}_1, w^A(\cdot)\big)$	$\texttt{Payoff}\big(\texttt{Pol}_1, w^B(\cdot)\big)$	...
$\texttt{Pol}_2 \in \mathcal{U}^{ad}$	$\texttt{Payoff}\big(\texttt{Pol}_2, w^A(\cdot)\big)$	$\texttt{Payoff}\big(\texttt{Pol}_2, w^B(\cdot)\big)$	...
...	...	...	...

In the robust or pessimistic approach, Nature is supposed to be malevolent, and the DM aims at protection against all odds

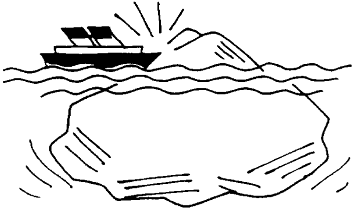

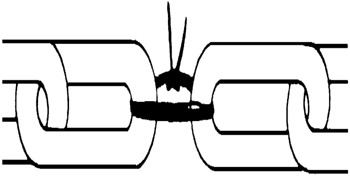

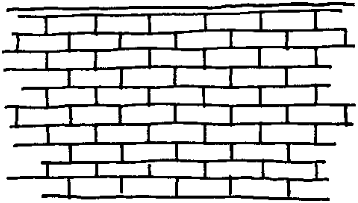

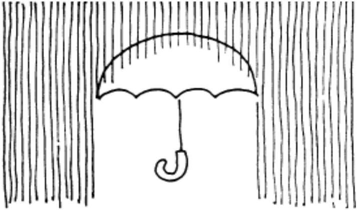

In the robust or pessimistic approach, Nature is supposed to be malevolent

- In the robust approach, the DM considers the worst payoff

$$\underbrace{\min_{w(\cdot)\in\Omega} \texttt{Payoff}\big(\texttt{Pol}, w(\cdot)\big)}_{\text{worst payoff}}$$

- Nature is supposed to be malevolent,
  and specifically selects the worst scenario:
  the DM plays after Nature has played, and maximizes the worst payoff

$$\max_{\texttt{Pol}\in\mathcal{U}^{ad}} \min_{w(\cdot)\in\Omega} \texttt{Payoff}\big(\texttt{Pol}, w(\cdot)\big)$$

- Robust, pessimistic, worst-case, maximin, minimax (for costs)

Guaranteed energy production

In a dam, the minimal energy production in a given period, corresponding to the worst water inflow scenario

The robust approach can be softened with plausibility weighting

- Let $\Theta : \Omega \to \mathbb{R} \cup \{-\infty\}$ be a a plausibility function.
- The higher, the more plausible:
  totally implausible scenarios are those for which $\Theta\big(w(\cdot)\big) = -\infty$
- Nature is malevolent, and specifically selects the worst scenario, but weighs it according to the plausibility function Θ
- The DM plays after Nature has played, and solves

$$\max_{\texttt{Pol} \in \mathfrak{U}^{ad}} \left[\min_{w(\cdot) \in \Omega} \left(\texttt{Payoff}\big(\texttt{Pol}, w(\cdot)\big) - \underbrace{\Theta\big(w(\cdot)\big)}_{\text{plausibility}} \right) \right]$$

In the optimistic approach, Nature is supposed to benevolent

Future. That period of time in which our affairs prosper, our friends are true and our happiness is assured.

Ambrose Bierce

- Instead of maximizing the worst payoff as in a robust approach, the optimistic focuses on the most favorable payoff

$$\underbrace{\max_{w(\cdot)\in\Omega} \texttt{Payoff}\big(\texttt{Pol}, w(\cdot)\big)}_{\text{best payoff}}$$

- Nature is supposed to benevolent, and specifically selects the best scenario: the DM plays after Nature has played, and solves

$$\max_{\texttt{Pol}\in\mathcal{U}^{ad}} \max_{w(\cdot)\in\Omega} \texttt{Payoff}\big(\texttt{Pol}, w(\cdot)\big)$$

The Hurwicz criterion reflects an intermediate attitude between optimistic and pessimistic approaches

A proportion $\alpha \in [0,1]$ graduates the level of prudence

$$\max_{\texttt{Pol} \in \mathcal{U}^{ad}} \left\{ \alpha \overbrace{\min_{w(\cdot) \in \Omega} \texttt{Payoff}\big(\texttt{Pol}, w(\cdot)\big)}^{\text{pessimistic}} + (1-\alpha) \underbrace{\max_{w(\cdot) \in \Omega} \texttt{Payoff}\big(\texttt{Pol}, w(\cdot)\big)}_{\text{optimistic}} \right\}$$

In the stochastic or expected approach, Nature is supposed to play stochastically

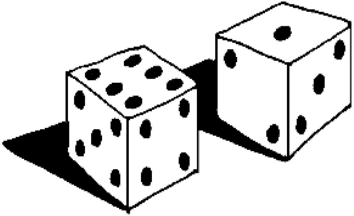

In the stochastic or expected approach, Nature is supposed to play stochastically

- The expected payoff is

$$\overbrace{\mathbb{E}\Big[\texttt{Payoff}\big(\texttt{Pol}, w(\cdot)\big)\Big]}^{\text{mean payoff}} = \sum_{w(\cdot)\in\Omega} \mathbb{P}\{w(\cdot)\}\texttt{Payoff}\big(\texttt{Pol}, w(\cdot)\big)$$

- Nature is supposed to play stochastically, according to distribution $\mathbb{P}$: the DM plays after Nature has played, and solves

$$\max_{\texttt{Pol}\in\mathfrak{U}^{ad}} \mathbb{E}\Big[\texttt{Payoff}\big(\texttt{Pol}, w(\cdot)\big)\Big]$$

- The discounted expected utility is the special case

$$\mathbb{E}\Big[\sum_{t=t_0}^{+\infty} \delta^{t-t_0} \mathtt{L}\big(x(t), u(t), w(t)\big)\Big]$$

The expected utility approach distorts payoffs before taking the expectation

- We consider a utility function $\mathtt{L}$ to assess the utility of the payoffs (for instance a CARA exponential utility function)
- The expected utility is

$$\underbrace{\mathbb{E}\left[\mathtt{L}\Big(\mathtt{Payoff}\big(\mathtt{Pol}, w(\cdot)\big)\Big)\right]}_{\text{expected utility}} = \sum_{w(\cdot)\in\Omega} \mathbb{P}\{w(\cdot)\}\mathtt{L}\Big(\mathtt{Payoff}\big(\mathtt{Pol}, w(\cdot)\big)\Big)$$

- The expected utility maximizer solves

$$\max_{\mathtt{Pol}\in\mathfrak{U}^{ad}} \mathbb{E}\left[\mathtt{L}\Big(\mathtt{Payoff}\big(\mathtt{Pol}, w(\cdot)\big)\Big)\right]$$

The ambiguity or multi-prior approach combines robust and expected criterion

- Different probabilities $\mathbb{P}$, termed as beliefs or priors and belonging to a set $\mathcal{P}$ of admissible probabilities on Ω
- The multi-prior approach combines robust and expected criterion by taking the worst beliefs in terms of expected payoff

$$\max_{\texttt{Pol}\in\mathfrak{U}^{ad}} \underbrace{\min_{\mathbb{P}\in\mathcal{P}} \overbrace{\mathbb{E}^{\mathbb{P}}\Big[\texttt{Payoff}\big(\texttt{Pol}, w(\cdot)\big)\Big]}^{\text{mean payoff}}}_{\text{pessimistic over probabilities}}$$

Convex risk measures cover a wide range of risk criteria

- Different probabilities $\mathbb{P}$, termed as beliefs or priors and belonging to a set $\mathcal{P}$ of admissible probabilities on Ω
- To each probability $\mathbb{P}$ is attached a plausibility $\Theta(\mathbb{P})$

$$\max_{\texttt{Pol} \in \mathfrak{U}^{ad}} \underbrace{\min_{\mathbb{P} \in \mathcal{P}} \overbrace{\mathbb{E}^{\mathbb{P}}\Big[\texttt{Payoff}\big(\texttt{Pol}, w(\cdot)\big)\Big]}^{\text{mean payoff}} - \overbrace{\Theta(\mathbb{P})}^{\text{plausibility}}}_{\text{pessimistic over probabilities}}$$

Non convex risk measures can lead to non diversification

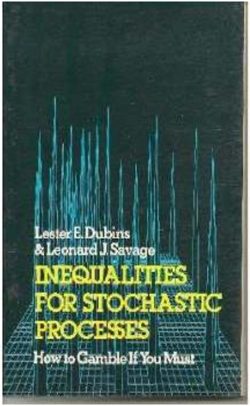

How to gamble if you must, L.E. Dubbins and L.J. Savage, 1965

Imagine yourself at a casino with \$1,000. For some reason, you desperately need \$10,000 by morning; anything less is worth nothing for your purpose.

The only thing possible is to gamble away your last cent, if need be, in an attempt to reach the target sum of \$10,000.

- The question is how to play, not whether. What ought you do? How should you play?
  - Diversify, by playing 1 \$ at a time?
  - Play boldly and concentrate, by playing 10,000 \$ only one time?
- What is your decision criterion?

Savage's minimal regret criterion... "Had I known"

$$\min_{\mathtt{Pol} \in \mathfrak{U}^{ad}} \Big\{ \overbrace{\max_{w(\cdot)\in\Omega} \underbrace{\Big[\max_{\text{anticipative policies } \overline{\mathtt{Pol}}} \mathtt{Payoff}\big(\overline{\mathtt{Pol}}, w(\cdot)\big) - \mathtt{Payoff}\big(\mathtt{Pol}, w(\cdot)\big) \Big]}_{\text{regret}}}^{\text{worst regret}} \Big\}$$

- ▷ If the DM knows the future in advance, she solves $\max_{\text{anticipative policies } \overline{\mathtt{Pol}}} \mathtt{Payoff}\big(\overline{\mathtt{Pol}}, w(\cdot)\big)$, for each scenario $w(\cdot) \in \Omega$
- ▷ The regret attached to a non-anticipative policy $\mathtt{Pol} \in \mathfrak{U}^{ad}$ is the loss due to not being visionary
- ▷ The best a non-visionary DM can do with respect to regret is minimizing it

Outline of the presentation

Upon which online information are decisions made?

We navigate between two stumbling blocks: rigidity and wizardry

- ▷ On the one hand, it is suboptimal to restrict oneself, as in the deterministic case, to open-loop controls depending only upon time, thereby ignoring the available information at the moment of making a decision
- ▷ On the other hand, it is impossible to suppose that we know in advance what will happen for all times: clairvoyance is impossible as well as look-ahead solutions

The in-between is **non-anticipativity constraint**

There are two ways to express the non-anticipativity constraint

Denote the uncertainties at time t by $w(t)$, and the control by $u(t)$

- ▷ Functional approach
  The control $u(t)$ may be looked after under the form

$$u(t) = \phi_t\big(\underbrace{w(t_0), \dots, w(t-1)}_{\text{past}}\big)$$

where ϕ_t is a function, called policy, strategy or decision rule

- ▷ Algebraic approach
  When uncertainties are considered as random variables (measurable mappings), the above formula for $u(t)$ expresses the measurability of the control variable $u(t)$ with respect to the past uncertainties, also written as

$$\underbrace{\sigma(u(t))}_{\sigma\text{-algebra}} \subset \sigma\big(\underbrace{w(t_0), \dots, w(t-1)}_{\text{past}}\big)$$

What is a solution at time t?

- In deterministic control, the solution $u(t)$ at time t is a single vector
- In stochastic control, the solution $u(t)$ at time t is a random variable expressed
  - either as $u(t) = \phi_t\big(w(t_0), \dots, w(t-1)\big)$, where $\phi_t : \mathbb{W}^{t-t_0} \to \mathbb{R}$
  - or as $u(t) : \Omega \to \mathbb{R}$ with measurability constraint $\sigma(u(t)) \subset \sigma\big(w(t_0), \dots, w(t-1)\big)$ or
    $u(t) = \mathbb{E}\left(u(t) \;\middle|\; w(t_0), \dots, w(t-1)\right)$
- Now, as time t goes on, the domain of the function ϕ_t expands, and so do the conditions $\sigma(u(t)) \subset \sigma\big(w(t_0), \dots, w(t-1)\big)$
- Therefore, for numerical reasons, the information $\big(w(t_0), \dots, w(t-1)\big)$ has to be compressed or approximated

Scenarios can be organized like a comb or like a tree

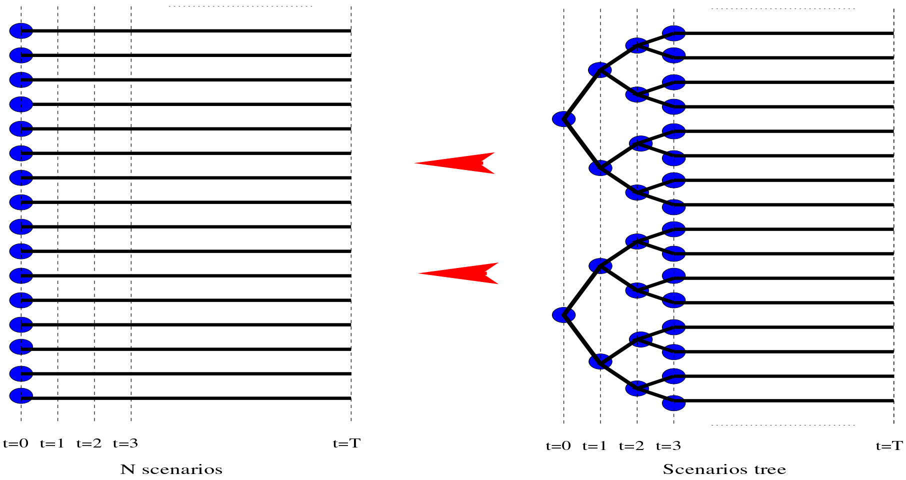

There are two classical ways to compress information

- **State-based functional approach**
  In the special case of the Markovian framework with $\big(w(t_0), \dots, w(T)\big)$ white noise, there is no loss of optimality to look for solutions as

$$u(t) = \psi_t \underbrace{\big(x(t)\big)}_{\text{state}} \quad \text{where} \quad \underbrace{x(t) \in \mathbb{X}}_{\text{fixed space}}, \quad \underbrace{x(t+1) = F_t\big(x(t), u(t), w(t)\big)}_{\text{dynamical equation}}$$

- **Scenario-based measurability approach**
  Scenarios are approximated by a finite family $\big(w^s(t_0), \dots, w^s(T)\big)$, $s \in S$
  - Either solutions $u^s(t)$ are indexed by $s \in S$ with the constraint that

$$\big(w^s(t_0), \dots, w^{s'}(t-1)\big) = \big(w^{s'}(t_0), \dots, w^{s'}(t-1)\big) \Rightarrow u^s(t) = u^{s'}(t)$$

  - Or — in the case of the scenario tree approach, where the scenarios $\big(w^s(t_0), \dots, w^s(T)\big)$, $s \in S$, are organized in a tree — solutions $u^n(t)$ are indexed by nodes n on the tree

More on what is a solution at time t State-based approach $u(t) = \psi_t(x(t))$

- The mapping ψ_t can be computed in advance (that is, at initial time t_0) and evaluated at time t on the available online information at that time t
  - either exactly (for example, by dynamic programming)
  - or approximately (for example, among linear decision rules) because the computational burden of finding *any* function is heavy
- The value $u(t) = \psi_t(x(t))$ can be computed at time t
  - either exactly by solving a proper optimization problem, which raises issues of dynamic consistency
  - or approximately (for example, by assuming that controls from time t on are open-loop)

More on what is a solution at time t Scenario-based approach

- An optimal "solution" can be computed scenario by scenario, with the problem that we obtain solutions such that

$$\big(w^s(t_0), \dots, w^s(t-1)\big) = \big(w^{s'}(t_0), \dots, w^{s'}(t-1)\big) \text{ and } u^s(t) \neq u^{s'}(t)$$

- Optimal solutions can be computed scenario by scenario and then merged (for example, by Progressive Hedging) to be forced to satisfy

$$\big(w^s(t_0), \dots, w^s(t-1)\big) = \big(w^{s'}(t_0), \dots, w^{s'}(t-1)\big) \Rightarrow u^s(t) = u^{s'}(t)$$

- The value $u(t)$ can be computed at time t depending on $\big(w^s(t_0), \dots, w^s(t-1)\big)$
  - either exactly by solving a proper optimization problem, which raises issues of dynamic consistency
  - or approximately (for example, by a sequence of two-stages problems)

Outline of the presentation

Where do we stand?

- ▷ How one frames the non-anticipativity constraint impacts numerical resolution methods
- ▷ On a finite scenario space, one obtains large (deterministic) optimization problems either on a tree or on a comb
- ▷ Else, one resorts to state-based formulations, with solutions as policies (dynamic programming)

Optimization approaches to attack complexity

Linear programming

- ▷ linear equations and inequalities
- ▷ no curse of dimension

Stochastic programming

- ▷ no special treatment of time and uncertainties
- ▷ no independence assumption
- ▷ decisions are indexed by a scenario tree
- ▷ what if information is not a node in the tree?

State-based dynamic optimization

- ▷ nonlinear equations and inequalities
- ▷ curse of dimensionality
- ▷ independence assumption on uncertainties
- ▷ special treatment of time (dynamic programming equation)
- ▷ decisions are indexed by an information state (feedback synthesis)
- ▷ an information state summarizes past controls and uncertainties
- ▷ decomposition-coordination methods to overcome the curse of dimensionality?

Summary

- *Stochastic* optimization highlights risk attitudes tackling
- Stochastic *dynamic* optimization emphasizes
  the handling of online information
- Many issues are raised, because
  - many ways to represent risk (criterion, constraints)
  - many information structures
  - tremendous numerical obstacles to overcome
- Each method has its numerical wall
  - in dynamic programming, the bottleneck is the dimension of the state (no more than 3)
  - in stochastic programming, the bottleneck is the number of stages (no more than 2)

Outline of the presentation

From linear to stochastic programming

- The linear program

$$\begin{aligned} \min \langle c \,, x \rangle & \\ x & \geq 0 \\ Ax + b & \geq 0 \end{aligned}$$

- becomes a stochastic program

$$\begin{aligned} \min \mathbb{E}(\langle c(\xi) \,, x \rangle) & \\ x & \geq 0 \\ A(\xi)x + b(\xi) & \geq 0 \end{aligned}$$

where $\xi : \Omega \to \Xi$ is a finite random variable

- so that there are as many inequalities as there are possible values for ξ

$$A\big(\xi(\omega)\big)x + b\big(\xi(\omega)\big) \geq 0 \,, \ \ \forall \omega \in \Omega$$

and these inequality constraints may define an empty domain for optimization

Recourse variables need be introduced for feasability issues

- We denote by $\xi \in \Xi$ any possible value of the random variable $\boldsymbol{\xi}$
- and we introduce a recourse variable $y = (y(\xi), \xi \in \Xi)$ and the program

$$\min \sum_{\xi \in \Xi} \mathbb{P}\{\xi\} \Big(\langle c(\xi) , x \rangle + \langle p(\xi) , y(\xi) \rangle \Big)$$

$$\begin{aligned} x &\geq 0 \\ y(\xi) &\geq 0 \, , \quad \forall \xi \in \Xi \\ A(\xi)x + b(\xi) - y(\xi) &\geq 0 \, , \quad \forall \xi \in \Xi \end{aligned}$$

- so that the inequality $A(\xi)x + b(\xi) - y(\xi) \geq 0$ is now possible, at (unitary recourse) price vector $p = (p(\xi), \xi \in \Xi)$
- As there are as many inequalities $A(\xi)x + b(\xi) - y(\xi) \geq 0$ as there are possible values for ξ, hence stochastic programs are huge problems, but can remain linear

Two-step stochastic programs with recourse can become deterministic non-smooth convex problems

- Define

$$Q(\xi, x) = \min\{\langle p(\xi) \, , y \rangle \, , A(\xi)x + b(\xi) - y \geq 0\}$$

which is a convex function of x, non-smooth

- so that the original two-step stochastic program with recourse

$$\min \sum_{\xi \in \Xi} \mathbb{P}\{\xi\} \langle c(\xi) \, , x \rangle + \langle p(\xi) \, , y(\xi) \rangle$$

$$\begin{aligned} x &\geq 0 \\ y(\xi) &\geq 0 \, , \quad \forall \xi \in \Xi \\ A(\xi)x + b(\xi) - y(\xi) &\geq 0 \, , \quad \forall \xi \in \Xi \end{aligned}$$

- now becomes the deterministic non-smooth convex problem

$$\min \langle c \, , x \rangle + \sum_{\xi \in \Xi} \mathbb{P}\{\xi\} Q(\xi, x)$$

$$x \geq 0$$

Roger Wets example

http://cermics.enpc.fr/~delara/ENSEIGNEMENT/

CEA-EDF-INRIA_2012/Roger_Wets1.pdf

Solutions of multi-stage stochastic optimization problems, without dual effect, can be indexed by a tree

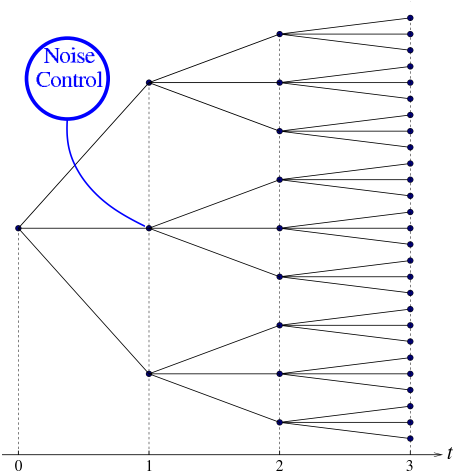

- ▷ Conditional probabilities given on the arcs, probabilities on the leafs
- ▷ Solutions indexed by the nodes of the tree

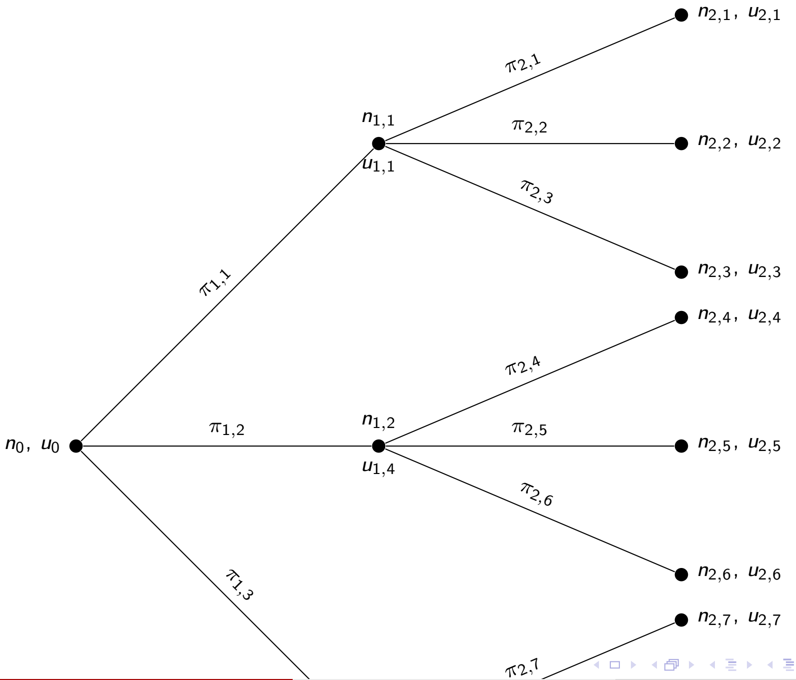
n2,1 , u2,1
π2,1
n1,1
π2,2
n2,2 , u2,2
u1,1
π2,3
π1,1
n2,3 , u2,3
n2,4 , u2,4
π2,4
n1,2
π2,5
π1,2
n0 , u0
n2,5 , u2,5
u1,4
π2,6
π1,3
n2,6 , u2,6
n2,7 , u2,7
π2,7

Outline of the presentation

Outline of the presentation

Outline of the presentation

Smart Power Systems, Renewable Energies and Markets: the Optimization Challenge

Michel DE LARA
CERMICS, École des Ponts ParisTech, France
and
Pierre Carpentier, ENSTA ParisTech, France
Jean-Philippe Chancelier, École des Ponts ParisTech, France
Pierre Girardeau, Artelys, France
Jean-Christophe Alais, Artelys, France
Vincent Leclère, École des Ponts ParisTech, France

CERMICS, France

November 14, 2014

During the night of 16 June 2013, electricity prices were negative

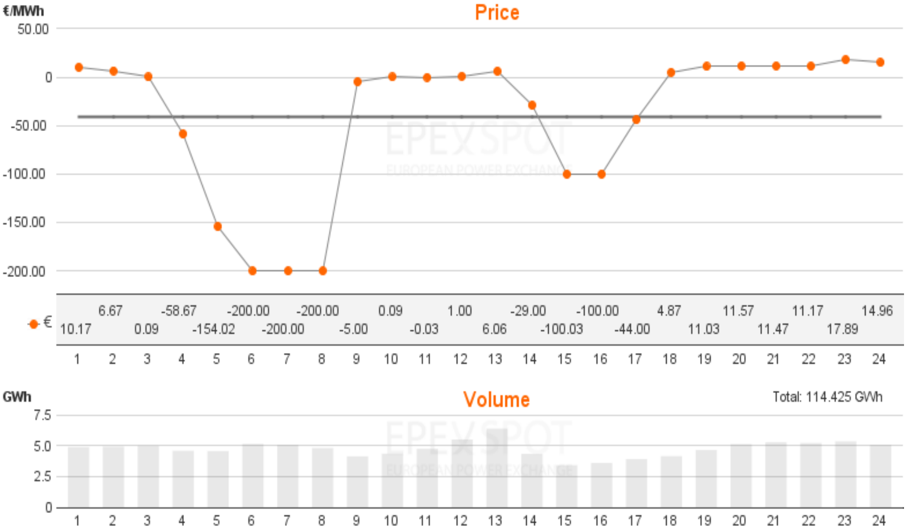

Outline of the talk

- In 2000, the *Optimization and Systems* team was created at École des Ponts ParisTech and, since then, we have *trained PhD students* in stochastic optimization, mostly with *Électricité de France Research and Development*
- Since 2011, we witness a growing demand from energy firms for stochastic optimization, fueled by a *deep and fast transformation of power systems*
- Renewable energies penetration, telecommunication technologies and markets remold power systems and challenge optimization
- More renewable energies $\rightarrow$ more unpredicability $+$ more variability $\rightarrow$
  - more storage $\rightarrow$ more dynamic optimization, optimal control
  - more stochastic optimization

  hence, stochastic optimal control
- We shed light on the two main *new issues* in *stochastic control* in comparison with *deterministic* control: *risk* attitudes and online *information*
- We cast a glow on two snapshots highlighting *ongoing research* in the field of stochastic control applied to energy

Outline of the presentation

1. Long term industry-academy cooperation
2. The remolding of power systems seen from an optimizer perspective
3. Moving from deterministic to stochastic dynamic optimization
4. Two snapshots on ongoing research
5. A need for training and research

Outline of the presentation

Outline of the presentation

1. Long term industry-academy cooperation
   - École des Ponts ParisTech–Cermics–Optimization and Systems
   - Industry partners of the Optimization and Systems Group
2. The remolding of power systems seen from an optimizer perspective
   - The remolding of power systems
   - Optimization is challenged
3. Moving from deterministic to stochastic dynamic optimization
   - Working out a toy example
   - Expliciting risk attitudes
   - Handling online information
   - Discussing framing and resolution methods
4. Two snapshots on ongoing research
   - Decomposition-coordination optimization methods under uncertainty
   - Risk constraints in optimization
5. A need for training and research

École des Ponts ParisTech is one of the world's oldest engineering institutes

- The École nationale des ponts et chaussées was founded in 1747 and is one of the world's oldest engineering institutes
- École des Ponts ParisTech is traditionally considered as belonging to the 5 leading engineering schools in France
- Young graduates find positions in professional sectors like transport and urban planning, banking, finance, consulting, civil works, industry, environnement, energy…
- Faculty and staff
  - 217 employees (including 50 subsidaries).
  - 165 module leaders, including 68 professors.
  - 1509 students.
- École des Ponts ParisTech is part of University Paris-Est

École des Ponts ParisTech hosts a substantial research activity

- Figures on research
  - Research personnel: 220
  - About 40 École des Ponts PhDs students graduate each year
- 10 research centers
  * CEREA (atmospheric environment), joint École des Ponts-EDF R&D
  * CEREVE (water, urban and environment studies)
  * CERMICS (mathematics and scientific computing)
  * CERTIS (information technologies and systems)
  * CIRED (international environment and development)
  * LATTS (techniques, regional planning and society)
  * LVMT (city, mobility, transport)
  * UR Navier (mechanics, materials and structures of civil engineering, geotechnic)
  * Saint-Venant laboratory (fluid mechanics), joint École des Ponts-EDF R&D
  * Paris School of Economic PSE

The CERMICS is the Centre d'enseignement et de recherche en mathématiques et calcul scientifique

- The scientific activity of CERMICS covers several domains in
  - scientific computing
  - modelling
  - optimization
- 15 senior researchers
  - 15 PhD
  - 12 habilitation à diriger des recherches
- Three missions
  - Teaching and PhD training
  - Scientific publications
  - Contracts
- 550 000 euros of contracts per year with
  - research and development centers of large industrial firms: CEA, CNES, EADS, EDF, Rio Tinto...
  - public research contracts

The Optimization and Systems Group comprises 3 senior researchers, as well as PhD students and external associated researchers

- Three senior researchers
  - J.-P. CHANCELIER
  - M. DE LARA
  - F. MEUNIER
- Eight PhD students
- Four associated researchers
  - P. CARPENTIER (ENSTA ParisTech)
  - L. ANDRIEU (EDF R&D)
  - K. BARTY (EDF R&D)
  - A. DALLAGI (EDF R&D)

Optimization and Systems Group research specialities

- Methods
  - Stochastic optimal control (discrete-time)
    - Large-scale systems
    - Discretization and numerical methods
    - Probability constraints
  - Discrete mathematics; combinatorial optimization
  - System control theory, viability and stochastic viability
  - Numerical methods for fixed points computation
  - Uncertainty and learning in economics
- Applications
  - Optimized management of power systems under uncertainty (production scheduling, power grid operations, risk management)
  - Transport modelling and management
  - Natural resources management (fisheries, mining, epidemiology)
- Softwares
  - Scicoslab, NSP
  - Oadlibsim

Publications since 2000

- ▷ 24 publications in peer-reviewed international journals
- ▷ 3 publications in collective works
- ▷ 4 books
  - ▹ Modeling and Simulation in Scilab/Scicos with ScicosLab 4.4 (2e édition, Springer-Verlag)
  - ▹ Introduction à SCILAB (2e édition, Springer-Verlag)
  - ▹ Sustainable Management of Natural Resources. Mathematical Models and Methods (Springer-Verlag)
  - ▹ Control Theory for Engineers (Springer-Verlag)
- ▷ 1 book submitted to Springer-Verlag
  - ▹ *Stochastic Optimization. At the Crossroads between Stochastic Control and Stochastic Programming*

Teaching

- Masters
  - *Master Parisien de Recherche Opérationnelle*
  - *Optimisation & Théorie des Jeux. Modélisation en Economie*
  - *Mathématiques, Informatique et Applications*
  - *Économie du Développement Durable, de l'Environnement et de l'Énergie*
  - *Renewable Energy Science and Technology Master ParisTech*
- École des Ponts ParisTech
  - Introduction à la recherche opérationnelle (F. MEUNIER)
  - Optimisation et contrôle (J.-P. CHANCELIER)
  - Modéliser l'aléa (J.-P. CHANCELIER)
  - Modélisation pour la gestion durable des ressources naturelles (M. DE LARA)

Industrial contracts mostly deal with energy issues, public ones touch on biodiversity management

- Industrial contracts
  - Conseil français de l'énergie (CFE)
  - SETEC Energy Solutions
  - Électricité de France (EDF R&D)
  - Thales
  - Institut français de l'énergie (IFE)
  - Gaz de France (GDF)
  - PSA
- Public contracts
  - STIC-AmSud (CNRS-INRIA-Affaires étrangères)
  - Centre d'étude des tunnels
  - CNRS ACI Écologie quantitative
  - RTP CNRS

Outline of the presentation

We cooperate with industry partners, looking for longlasting research relations through training and capacity building

- As academics, we cooperate with industry partners, looking for longlasting close relations
- We are not consultants working for clients, but focus en capacity building
- Our job consists mainly in
  - training Master and PhD students, working within the company and interacting with us, on subjects designed jointly
  - developing methods, algorithms
  - contributing to computer codes developed within the company
  - training professional engineers in the company

Électricité de France R & D / Département OSIRIS

- Électricité de France is the French electricity main producer
  - 159 000 collaborateurs dans le monde
  - 37 millions de clients dans le monde
  - 65,2 milliards d'euros de chiffre d'affaire
  - 630,4 TWh produits dans le monde
- Électricité de France Research & Development
  - 486 millions d'euros de budget
  - 2 000 personnes
- Département OSIRIS
  Optimisation, simulation, risques et statistiques pour les marchés de l'énergie
  Optimization, simulation, risks and statistics for the energy markets
  - 145 salariés (dont 10 doctorants)
  - 25 millions d'euros de budget

What is “optimization”?

Optimizing is obtaining the best compromise between needs and resources

Marcel Boiteux (président d’honneur d’Électricité de France)

- Resources: portfolio of assets
  - production units
    - costly/not costly: thermal/hydropower
    - stock/flow, predictable/unpredictable: thermal/wind
  - tariffs options, contracts
- Needs: energy, safety, environment
  - energy uses
  - safety, quality, resilience (breakdowns, blackout)
  - environment protection (pollution) and alternative uses (dam water)
- Best compromise: minimize socio-economic costs (including externalities)

The Optimization and Systems Group has trained 10 PhD from 2004 to 2014, most of them related with EDF and energy management

* Laetitia ANDRIEU, former PhD student at EDF, now researcher EDF
* Kengy BARTY, former PhD student at EDF, now researcher EDF
* Daniel CHEMLA, former PhD student
* Anes DALLAGI, former PhD student at EDF, now researcher EDF
* Laurent GILOTTE, former PhD student with IFE, researcher EDF
* Pierre GIRARDEAU, former PhD student at EDF, now with ARTELYS
* Eugénie LIORIS, former PhD student
* Babacar SECK, former PhD student at EDF
* Cyrille STRUGAREK, former PhD student at EDF, now with Munich-Ré
* Jean-Christophe ALAIS, former PhD student at EDF, now with ARTELYS
* Vincent LECLERE, former PhD student (partly at EDF), now with CERMICS

PhD subjects reflect academic issues raised by industrial problems

- *Contributions to the Discretization of Measurability Constraints for Stochastic Optimization Problems*,
- *Optimization under Probability Constraint*,
- *Uncertainty, Inertia and Optimal Decision. Optimal Control Models Applied to Greenhouse Gas Abatment Policies Selection*,
- *Variational Approaches and other Contributions in Stochastic Optimization*,
- *Particular Methods in Stochastic Optimal Control*,
- *From Risk Constraints in Stochastic Optimization Problems to Utility Functions*,
- *Resolution of Large Size Problems in Dynamic Stochastic Optimization and Synthesis of Control Laws*,
- *Evaluation and Optimization of Collective Taxis Systems*,
- *Risk and Optimization for Energies Management*,
- *Risk, Optimization, Large Systems*,

Recently, contacts have expanded with small companies

- ARTELYS is a company specializing in optimization, decision-making and modeling. Relying on their high level of expertise in quantitative methods, the consultants deliver efficient solutions to complex business problems. They provide services to diversified industries: Energy & Environment, Logistics & Transportation, Telecommunications, Finance and Defense.
- Créée en 2011, SETEC Energy Solutions est la filiale du groupe SETEC spécialisée dans les domaines de la production et de la maîtrise de l'énergie en France et à l'étranger. SETEC Energy Solutions apporte à ses clients la maîtrise des principaux process énergétiques pour la mise en œuvre de solutions innovantes depuis les phases initiales de définition d'un projet jusqu'à son exploitation.
- SUN'R Smart Energy is a Paris based company with a focus on building smarter solutions for distributed energy resources in the context of emerging deregulated energy markets and a solid political will towards the development of both renewables and energy storage. The company is part of a larger group founded in 2007 and is a growing, well-funded early stage business.

French Energy Council, member of the World Energy Council, contracted the Optimization and Systems group to report on Optimization methods for smart grids

- Formed in 1923, the World Energy Council (WEC) is the UN-accredited global energy body, representing the entire energy spectrum, with more than 3000 member organisations located in over 90 countries and drawn from governments, private and state corporations, academia, NGOs and energy-related stakeholders
- WEC informs global, regional and national energy strategies by hosting high-level events, publishing authoritative studies, and working through its extensive member network to facilitate the world's energy policy dialogue
- In 2012, the French Energy Council contracted the Optimization and Systems group to produce a report on Optimization methods for smart grids

Summary

The following slides on the remolding of power systems express a viewpoint

- ▷ from an optimizer perspective
- ▷ working in an applied mathematics research center
- ▷ in a French engineering institute
- ▷ having contributed to train students now working in energy
- ▷ having contacts and contracts with energy/environment firms

Outline of the presentation

Outline of the presentation

Three key drivers are remolding power systems

- ▷ Environment
- ▷ Markets
- ▷ Technology

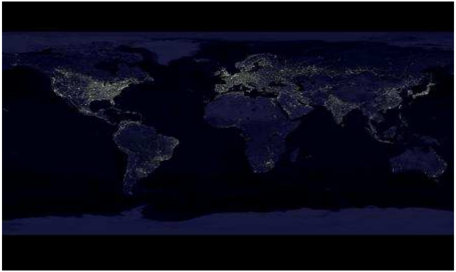

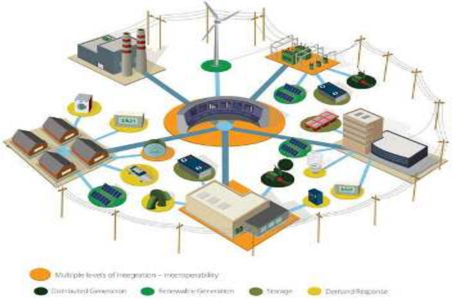

Key driver: environmental concern

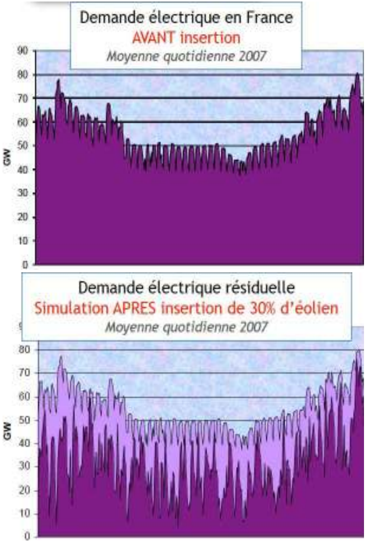

The European Union climate and energy package materializes an environmental concern with three 20-20-20 objectives for 2020

- ▷ a 20% improvement in the EU's energy efficiency
- ▷ a 20% reduction in EU greenhouse gas emissions from 1990 levels
- ▷ raising the share of EU energy consumption produced from renewable resources to 20%

⇓

Successfully integrating renewable energy sources has become critical, and made especially difficult because they are unpredictable and highly variable, hence triggering the use of local storage

Key driver: economic deregulation

- A power system (generation/transmission/distribution)
  - less and less vertical (deregulation of energy markets)
  - hence with many players with their own goals
- with some new players
  - industry (electric vehicle)
  - regional public authorities (autonomy, efficiency)
- with a network in horizontal expansion
  (the Pan European electricity transmission system counts 10,000 buses, 15,000 power lines, 2,500 transformers, 3,000 generators, 5,000 loads)
- with more and more exchanges (trade of commodities)

$$\Downarrow$$

A change of paradigm for management
from centralized to more and more decentralized

Key driver: telecommunication technology

Linky

A power system with more and more technology due to evolutions in the fields of metering, computing and telecoms

- ▷ smart meters
- ▷ sensors
- ▷ controllers
- ▷ grid communication devices. . .

$\Downarrow$

A huge amount of data which, one day, will be a new potential for optimized management

The "smart grid"? An infrastructure project with promises to be fulfilled by a "smart power system"

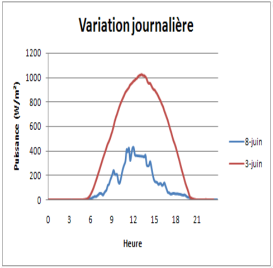

- Hardware / infrastructures / smart technologies
  - Renewable energies technologies
  - Smart metering
  - Storage
- Promises
  - Quality, tariffs
  - More safety
  - More renewables (environmentally friendly)
- Software / smart management
  (energy supply being less flexible,
  make the demand more flexible)

smart management, smart operation, smart meter management, smart distributed generation, load management, advanced distribution management systems, active demand management, diffuse effacement, distribution management systems, storage management, smart home, demand side management...

Outline of the presentation

1. Long term industry-academy cooperation
   - École des Ponts ParisTech–Cermics–Optimization and Systems
   - Industry partners of the Optimization and Systems Group
2. The remolding of power systems seen from an optimizer perspective
   - The remolding of power systems
   - Optimization is challenged
3. Moving from deterministic to stochastic dynamic optimization
   - Working out a toy example
   - Expliciting risk attitudes
   - Handling online information
   - Discussing framing and resolution methods
4. Two snapshots on ongoing research
   - Decomposition-coordination optimization methods under uncertainty
   - Risk constraints in optimization
5. A need for training and research

We witnessed a call from EDF to optimizers

- Every three years, Électricité de France (EDF) organizes an international *Conference on Optimization and Practices in Industry* (COPI)
- At the last COPI'11, Jean-François Faugeras from EDF R&D opened the conference with a plenary talk entitled "Smart grids: a wind of change in power systems and new opportunities for optimization"
- He claimed that "power system players are facing high level problems to solve requiring new optimization methods and tools",
  with "not only a 'smart(er)' grid but a 'smart(er)' power system"
  and called on the optimizers to develop new methods
- In 2012, EDF R&D has sponsored a new program *Gaspard Monge pour l'Optimisation et la recherche opérationnelle* (PGMO)
  to support academic research in the field of optimization

What is “optimization”?

Optimizing is obtaining the best compromise between needs and resources

Marcel Boiteux (président d’honneur d’Électricité de France)

- Resources: portfolio of assets
  - production units
    - costly/not costly: thermal/hydropower
    - stock/flow, predictable/unpredictable: thermal/wind
  - tariffs options, contracts
- Needs: energy, safety, environment
  - energy uses
  - safety, quality, resilience (breakdowns, blackout)
  - environment protection (pollution) and alternative uses (dam water)
- Best compromise: minimize socio-economic costs (including externalities)

Electrical engineers metiers and skills are evolving

- Unit commitment, optimal dispatch of generating units:
  finding the least-cost dispatch of available generation resources to meet the electrical load
  - which unit? 0/1 variables
  - which power level? continuous variables

  subject to more unpredictable energy flows (solar, wind)
  and demand (electrical devices, cars)
- Markets: day-ahead, intra-day (balancing market):
  dispatcher takes bids from the generators, demand forecasts from the distribution companies and clears the market
  subject to more unpredictability, more players
- Long term planning
  subject to more unpredictability (technologies, climate), more players
- . . . Without even speaking of voltage, frequency and phase control

Let us have a look at economic dispatch (static) as a cost-minimization problem under supply-demand balance

Consider energy production units $i = 1, \ldots, N$, like coal, gas, nuclear...

$$\underbrace{\min_{(u_1,\ldots,u_N)} \sum_{i=1}^{N} J_i(u_i)}_{\text{costs minimization}} \quad \text{under} \quad \underbrace{\sum_{i=1}^{N} \Theta_i(u_i) = D}_{\text{supply = demand}}$$

where

- u_i is the decision (production level) made for each unit i
- $J_i(u_i)$ is the cost of making decision u_i for unit i
- $\Theta_i(u_i)$ is the production induced by making decision u_i for unit i
- D is the demand

Inviting in uncertainty gives economic dispatch new suits of clothes

$$\underbrace{\min_{(u_1,\ldots,u_N)} \mathbb{E}\big[\sum_{i=1}^{N} J_i(u_i, \overbrace{p_i}^{\text{price}})\big]}_{\text{expected costs minimization}} \quad \text{under} \quad \underbrace{\sum_{i=1}^{N} \Theta_i(u_i, \overbrace{w_i}^{\text{weather}}) = \overbrace{D}^{\text{demand}}}_{\text{almost-surely supply = demand}}$$

- ▷ Mathematical description of sources of uncertainties (prices p_i, weather w_i, demand D, failures...): statistics? bounds?
- ▷ Mathematical formulation of the criterion under uncertainty: in expectation ($\mathbb{E}$)? worst case (max)?
- ▷ Mathematical formulation of the constraints under uncertainty: in expectation? in probability? almost surely? robust? by penalization?

With uncertainty come stages, hence a dynamics

- In electricity, the *supply matches demand* equation "is like gravity, you cannot negociate" (who claimed that?)
- One way or another, we are driven to add a new instantaneous source u_{N+1}

$$\sum_{i=1}^{N} \Theta_i(u_i, \overbrace{w_i}^{\text{weather}}) + \underbrace{u_{N+1}}_{\text{new source}} = \overbrace{D}^{\text{demand}}$$

- The control $u_{N+1} = D - \sum_{i=1}^{N} \Theta_i(u_i, w_i)$ depends on the uncertain variables D and $w_1, \ldots, w_N$
- Whereas $u_1, \ldots, u_N$ are decisions made before knowing their realizations
- To cut to the point, we now have two stages

$$\Downarrow$$

Piecing things together, we started from static economic dispatch and,
on the path of making allowance for uncertainty,
we have been quite naturally led to
dynamic economic dispatch under risk

Optimization skills will follow the power system evolution

We focus on generation and trading, not on transmission and distribution

- Less base production and more wind and photovoltaic fatal generation makes supply more unpredictible
  → stochastic optimization
- Hence more storage (batteries, pumping stations)
  → dynamical optimization, reserves dimensioning
- The shape of the load is changing due to electric vehicle penetration
  → demand-side management, "peak shaving", adaptive tariffs
- New subsystems emerge with local information and means of action: smart meters, new producers, micro-grid, virtual power plant
  → agregation, coordination, decentralized optimization
- Markets (day-ahead, intra-day)
  → optimization under uncertainty
- Environmental constraints on production (CO2) and resources usages (water)
  → risk constraints

Summary

- ▷ Three major key factors — environmental concern, deregulation, telecommunication, metering and computing technology — drive the power systems remolding
- ▷ This remolding induces a change of paradigm for management: from vertical centralized predictible "stock" energies to more horizontal decentralized unpredictible variable "flow" energies
- ▷ Specific optimization skills will be required, because an optimal solution is *balancing on a knife edge*, hence might perform poorly under off-nominal conditions, like a *too much adjusted suit cracking at the first move*

Roger Wets' illuminating example: deterministic vs. robust

of a furniture manufacturer deciding how many dressers of each of 4 types to produce, with carpentry and finishing man-hours as constraints; when the ten parameters become random, the stochastic optimal solution considers all $\approx 10^6$ possibilities and provides a robust solution $(257\,;0\,;665\,;34)$, whereas the deterministic solution $(1,333\,;0\,;0\,;67)$ does not point in the right direction

Outline of the presentation

We distinguish two polar classes of models: knowledge models *versus* decision models

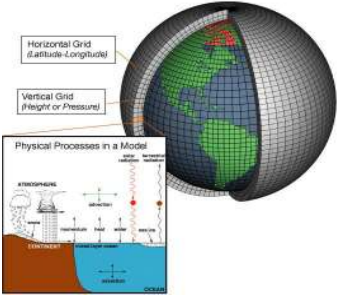

Knowledge models:
$1/1\ 000\ 000 \rightarrow 1/1\ 000 \rightarrow 1/1$ maps

Office of Oceanic and Atmospheric Research (OAR) climate model

We distinguish two polar classes of models: knowledge models *versus* decision models

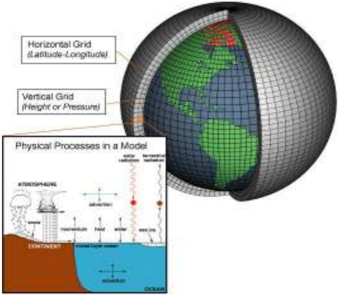

Knowledge models:
1/1 000 000 → 1/1 000 → 1/1 maps

Office of Oceanic and Atmospheric Research (OAR) climate model

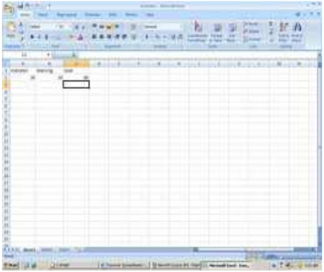

Action/decision models:
economic models are fables designed to provide insight

William Nordhaus economic-climate model

Outline of the presentation

1. Long term industry-academy cooperation
   - École des Ponts ParisTech–Cermics–Optimization and Systems
   - Industry partners of the Optimization and Systems Group
2. The remolding of power systems seen from an optimizer perspective
   - The remolding of power systems
   - Optimization is challenged
3. Moving from deterministic to stochastic dynamic optimization
   - Working out a toy example
   - Expliciting risk attitudes
   - Handling online information
   - Discussing framing and resolution methods
4. Two snapshots on ongoing research
   - Decomposition-coordination optimization methods under uncertainty
   - Risk constraints in optimization
5. A need for training and research

Let us work out a toy example of economic dispatch as a cost-minimization problem under supply-demand balance

- Production: consider two energy production units
  - a "cheap" limited one with which we can produce quantity q_0, with $0 \leq q_0 \leq q_0^\sharp$, at cost $c_0 q_0$
  - an "expensive" unlimited one with which we can produce quantity q_1, with $0 \leq q_1$, at cost $c_1 q_1$, with $c_1 > c_0$
- Consumption: the demand is $D \geq 0$
- Balance: ensuring at least the demand

$$D \leq q_0 + q_1$$

- Optimization: total costs minimization

$$\min_{q_0, q_1} \underbrace{c_0 q_0 + c_1 q_1}_{\text{total costs}}$$

When the demand D is deterministic, the optimization problem is well posed

- The deterministic demand D is a single number, and we consider

$$\min_{q_0, q_1} c_0 q_0 + c_1 q_1$$

$$\text{under the constraints} \quad \begin{array}{rl} 0 & \leq q_0 \leq q_0^\sharp \\ 0 & \leq q_1 \\ D & \leq q_0 + q_1 \end{array}$$

- The solution is $q_0^\star = \min\{q_0^\sharp, D\}$, $q_1^\star = [D - q_0^\sharp]_+$, that is,
  - if the demand D is below the capacity $q_0^\sharp$ of the "cheap" energy source

$$D \leq q_0^\sharp \Rightarrow q_0^\star = D \,, \quad q_1^\star = 0$$

  - if the demand D is above the capacity $q_0^\sharp$ of the "cheap" energy source,

$$D > q_0^\sharp \Rightarrow q_0^\star = q_0^\sharp \,, \quad q_1^\star = D - q_0^\sharp$$

- Now, what happens when the demand D is no longer deterministic?

If we know the demand beforehand, the optimization problem is deterministic

- We suppose that the demand is a random variable $D : \Omega \to \mathbb{R}_+$
- If we solve the problem for each possible value $D(\omega)$ of the random variable D, when $\omega \in \Omega$, we obtain

$$q_0(\omega) = \min\{q_0^\sharp, D(\omega)\} , \quad q_1(\omega) = [D(\omega) - q_0^\sharp]_+$$

  and we face an informational issue
- Indeed, we treat the demand D as if observed before making the decisions q_0 and q_1
- When the demand D is not observed, how can we do?

What happens if we replace the uncertain value D of the demand by its mean $\overline{D}$ in the deterministic solution?

- If we suppose that the demand D is a random variable $D : \Omega \to \mathbb{R}_+$, with mathematical expectation $\mathbb{E}(D) = \overline{D}$
- and that we propose the "deterministic solution"

$$q_0^{(\overline{D})} = \min\{q_0^\sharp, \overline{D}\} \ , \ \ q_1^{(\overline{D})} = [\overline{D} - q_0^\sharp]_+$$

- we cannot assure the inequality

$$\underbrace{D(\omega)}_{\text{uncertain}} \leq \underbrace{q_0 + q_1}_{\text{deterministic}} \ , \ \ \forall \omega \in \Omega$$

because $\sup_{\omega \in \Omega} D(\omega) > \overline{D} = q_0^{(\overline{D})} + q_1^{(\overline{D})}$

- Are there better solutions among the deterministic ones?

When the demand D is bounded above, the robust optimization problem has a solution

- In the robust optimization problem, we minimize

$$\min_{q_0, q_1} c_0 q_0 + c_1 q_1$$

$$\text{under the constraints} \quad \begin{aligned} 0 &\leq q_0 \leq q_0^{\sharp} \\ 0 &\leq q_1 \\ D(\omega) &\leq q_0 + q_1 \,, \quad \forall \omega \in \Omega \end{aligned}$$

- When $D^{\sharp} = \sup_{\omega \in \Omega} D(\omega) < +\infty$, the solution is $q_0^{\star} = \min\{q_0^{\sharp}, D^{\sharp}\} \,, \quad q_1^{\star} = [D^{\sharp} - q_0^{\sharp}]_+$
- Now, the total cost $c_0 q_0^{\star} + c_1 q_1^{\star}$ is an increasing function of the upper bound $D^{\sharp}$ of the demand
- Is it not too costly to optimize under the worst-case situation?

Where do we stand?

- ▷ When the demand D is deterministic, the optimization problem is well posed
- ▷ If we know the demand beforehand, the optimization problem is deterministic
- ▷ If we replace the uncertain value D of the demand by its mean $\overline{D}$ in the deterministic solution, we remain with a feasability issue
- ▷ When the demand D is bounded above, the robust optimization problem has a solution, but it is costly

Where do we stand?

- When the demand D is deterministic, the optimization problem is well posed
- If we know the demand beforehand, the optimization problem is deterministic
- If we replace the uncertain value D of the demand by its mean $\overline{D}$ in the deterministic solution, we remain with a feasability issue
- When the demand D is bounded above, the robust optimization problem has a solution, but it is costly

To overcome the above difficulties, we propose to introduce stages

$$\underbrace{D(\omega)}_{\text{uncertain}} \leq \underbrace{q_0}_{\text{deterministic}} + \underbrace{q_1(\omega)}_{\text{uncertain}} , \quad \forall \omega \in \Omega$$

- the decision q_0 is made before observing the demand $D(\omega)$
- the decision $q_1(\omega)$ is made after observing the demand $D(\omega)$

To overcome the above difficulties, we turn to stochastic optimization

- ▷ We suppose that the demand D is a random variable, and minimize

$$\min_{q_0, q_1} \mathbb{E}[c_0 q_0 + c_1 q_1]$$

$$\text{under the constraints} \quad \begin{array}{ll} 0 & \leq q_0 \leq q_0^\sharp \\ 0 & \leq q_1 \\ D & \leq q_0 + q_1 \\ q_1 & \text{depends upon } D \end{array}$$

and we emphasize two issues, new with respect to the deterministic case

- ▷ expliciting online information issue:
  the decision q_1 depends upon the random variable D
- ▷ expliciting risk attitudes:
  we aggregate the total costs with respect to all possible values
  by taking the expectation $\mathbb{E}[c_0 q_0 + c_1 q_1]$

Turning to stochastic optimization forces one to specify online information

- We suppose that the demand D is a random variable, and minimize

$$\min_{q_0, q_1} \mathbb{E}[c_0 q_0 + c_1 q_1]$$

$$\text{under the constraints} \quad \begin{array}{rl} 0 & \leq q_0 \leq q_0^{\sharp} \\ 0 & \leq q_1 \\ D & \leq q_0 + q_1 \\ q_1 & \text{depends upon } D \end{array}$$

- specifying that the decision q_1 depends upon the random variable D, whereas q_0 does not, forces to consider two stages and a so-called non-anticipativity constraint (more on that later)
  - first stage: q_0 does not depend upon the random variable D
  - second stage: q_1 depends upon the random variable D

Turning to stochastic optimization forces one to specify risk attitudes

▷ We suppose that the demand D is a random variable, and minimize

$$\min_{q_0, q_1} \mathbb{E}[c_0 q_0 + c_1 q_1]$$

$$\text{under the constraints} \quad \begin{array}{ll} 0 & \leq q_0 \leq q_0^\sharp \\ 0 & \leq q_1 \\ D & \leq q_0 + q_1 \\ q_1 & \text{depends upon } D \end{array}$$

▷ Now that q_1 depends upon the random variable D, it is also a random variable, and so is the total cost $c_0 q_0 + c_1 q_1$; therefore, we have to aggregate the total costs with respect to all possible values, and we chose to do it by taking the expectation $\mathbb{E}[c_0 q_0 + c_1 q_1]$

In the uncertain framework, two additional questions must be answered with respect to the deterministic case

Question (expliciting risk attitudes)

How are the uncertainties taken into account
in the payoff criterion and in the constraints?

Question (expliciting available online information)

Upon which online information are decisions made?

Outline of the presentation

The output of a stochastic optimization problem is a random variable. How can we rank random variables?

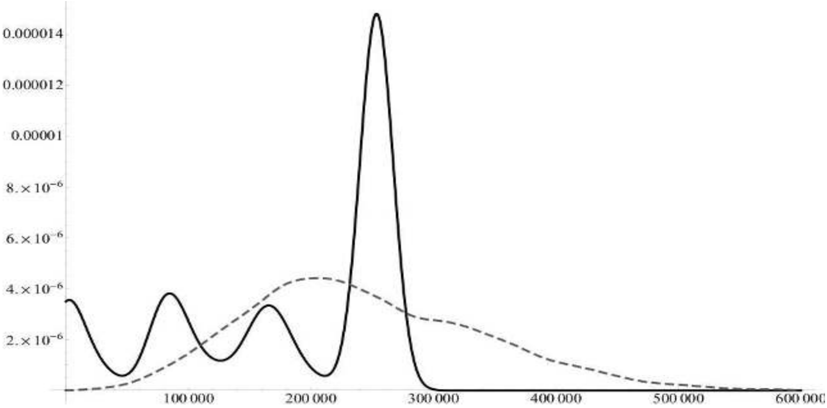

How are the uncertainties taken into account in the payoff criterion and in the constraints?

In a probabilistic setting, where uncertainties are random variables, a classical answer is

- ▷ to take the mathematical expectation of the payoff (risk-neutral approach)

$$\mathbb{E}(\text{payoff})$$

- ▷ and to satisfy all (physical) constrainsts almost surely that is, practically, for all possible issues of the uncertainties (robust approach)

$$\mathbb{P}(\text{constrainsts}) = 1$$

But there are many other ways to handle risk: robust, worst case, risk measures, in probability, almost surely, by penalization, etc.

A policy and a criterion yield a real-valued payoff

Given a policy $\texttt{Pol} \in \mathcal{U}^{ad}$ and a scenario $w(\cdot) \in \Omega$, we obtain a payoff

$$\texttt{Payoff}\big(\texttt{Pol}, w(\cdot)\big)$$

hence a mapping $\mathcal{U}^{ad} \times \Omega \to \mathbb{R}$

Policies/Scenarios	$w^A(\cdot) \in \Omega$	$w^B(\cdot) \in \Omega$	...
$\texttt{Pol}_1 \in \mathcal{U}^{ad}$	$\texttt{Payoff}\big(\texttt{Pol}_1, w^A(\cdot)\big)$	$\texttt{Payoff}\big(\texttt{Pol}_1, w^B(\cdot)\big)$	...
$\texttt{Pol}_2 \in \mathcal{U}^{ad}$	$\texttt{Payoff}\big(\texttt{Pol}_2, w^A(\cdot)\big)$	$\texttt{Payoff}\big(\texttt{Pol}_2, w^B(\cdot)\big)$	...
...	...	...	...

In the robust or pessimistic approach, Nature is supposed to be malevolent, and the DM aims at protection against all odds

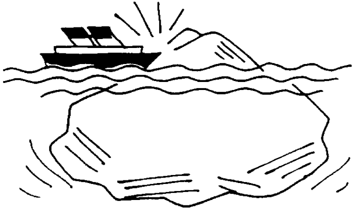
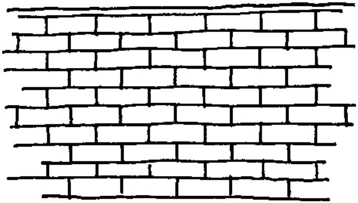
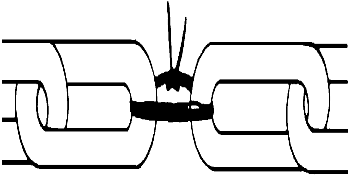
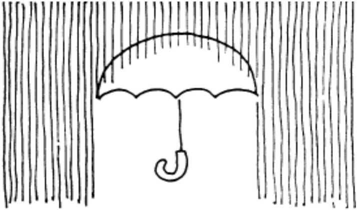

In the robust or pessimistic approach, Nature is supposed to be malevolent

- In the robust approach, the DM considers the worst payoff

$$\underbrace{\min_{w(\cdot)\in\Omega} \texttt{Payoff}\big(\texttt{Pol}, w(\cdot)\big)}_{\text{worst payoff}}$$

- Nature is supposed to be malevolent,
  and specifically selects the worst scenario:
  the DM plays after Nature has played, and maximizes the worst payoff

$$\max_{\texttt{Pol}\in\mathfrak{U}^{ad}} \min_{w(\cdot)\in\Omega} \texttt{Payoff}\big(\texttt{Pol}, w(\cdot)\big)$$

- Robust, pessimistic, worst-case, maximin, minimax (for costs)

Guaranteed energy production

In a dam, the minimal energy production in a given period, corresponding to the worst water inflow scenario

The robust approach can be softened with plausibility weighting

- Let $\Theta : \Omega \to \mathbb{R} \cup \{-\infty\}$ be a a plausibility function.
- The higher, the more plausible:
  totally implausible scenarios are those for which $\Theta\big(w(\cdot)\big) = -\infty$
- Nature is malevolent, and specifically selects the worst scenario, but weighs it according to the plausibility function Θ
- The DM plays after Nature has played, and solves

$$\max_{\texttt{Pol} \in \mathfrak{U}^{ad}} \left[\min_{w(\cdot) \in \Omega} \left(\texttt{Payoff}\big(\texttt{Pol}, w(\cdot)\big) - \underbrace{\Theta\big(w(\cdot)\big)}_{\text{plausibility}} \right) \right]$$

In the optimistic approach, Nature is supposed to benevolent

Future. That period of time in which our affairs prosper, our friends are true and our happiness is assured.

Ambrose Bierce

- Instead of maximizing the worst payoff as in a robust approach, the optimistic focuses on the most favorable payoff

$$\underbrace{\max_{w(\cdot)\in\Omega} \texttt{Payoff}\big(\texttt{Pol}, w(\cdot)\big)}_{\text{best payoff}}$$

- Nature is supposed to benevolent, and specifically selects the best scenario: the DM plays after Nature has played, and solves

$$\max_{\texttt{Pol}\in\mathfrak{U}^{ad}} \max_{w(\cdot)\in\Omega} \texttt{Payoff}\big(\texttt{Pol}, w(\cdot)\big)$$

The Hurwicz criterion reflects an intermediate attitude between optimistic and pessimistic approaches

A proportion $\alpha \in [0,1]$ graduates the level of prudence

$$\max_{\texttt{Pol} \in \mathcal{U}^{ad}} \left\{ \alpha \overbrace{\min_{w(\cdot) \in \Omega} \texttt{Payoff}\big(\texttt{Pol}, w(\cdot)\big)}^{\text{pessimistic}} + (1-\alpha) \underbrace{\max_{w(\cdot) \in \Omega} \texttt{Payoff}\big(\texttt{Pol}, w(\cdot)\big)}_{\text{optimistic}} \right\}$$

In the stochastic or expected approach, Nature is supposed to play stochastically

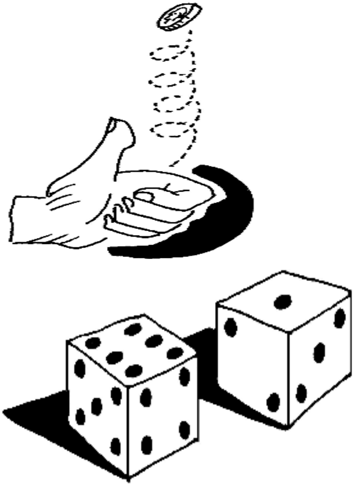

In the stochastic or expected approach, Nature is supposed to play stochastically

- The expected payoff is

$$\overbrace{\mathbb{E}\Big[\mathtt{Payoff}\big(\mathtt{Pol}, w(\cdot)\big)\Big]}^{\text{mean payoff}} = \sum_{w(\cdot)\in\Omega} \mathbb{P}\{w(\cdot)\}\mathtt{Payoff}\big(\mathtt{Pol}, w(\cdot)\big)$$

- Nature is supposed to play stochastically, according to distribution $\mathbb{P}$: the DM plays after Nature has played, and solves

$$\max_{\mathtt{Pol}\in\mathfrak{U}^{ad}} \mathbb{E}\Big[\mathtt{Payoff}\big(\mathtt{Pol}, w(\cdot)\big)\Big]$$

- The discounted expected utility is the special case

$$\mathbb{E}\Big[\sum_{t=t_0}^{+\infty} \delta^{t-t_0} \mathtt{L}\big(x(t), u(t), w(t)\big)\Big]$$

The expected utility approach distorts payoffs before taking the expectation

- We consider a utility function $\mathtt{L}$ to assess the utility of the payoffs (for instance a CARA exponential utility function)
- The expected utility is

$$\underbrace{\mathbb{E}\left[\mathtt{L}\Big(\mathtt{Payoff}\big(\mathtt{Pol}, w(\cdot)\big)\Big)\right]}_{\text{expected utility}} = \sum_{w(\cdot)\in\Omega} \mathbb{P}\{w(\cdot)\}\mathtt{L}\Big(\mathtt{Payoff}\big(\mathtt{Pol}, w(\cdot)\big)\Big)$$

- The expected utility maximizer solves

$$\max_{\mathtt{Pol}\in\mathfrak{U}^{ad}} \mathbb{E}\left[\mathtt{L}\Big(\mathtt{Payoff}\big(\mathtt{Pol}, w(\cdot)\big)\Big)\right]$$

The ambiguity or multi-prior approach combines robust and expected criterion

- Different probabilities $\mathbb{P}$, termed as beliefs or priors and belonging to a set $\mathcal{P}$ of admissible probabilities on Ω
- The multi-prior approach combines robust and expected criterion by taking the worst beliefs in terms of expected payoff

$$\max_{\texttt{Pol} \in \mathfrak{U}^{ad}} \underbrace{\min_{\mathbb{P} \in \mathcal{P}} \overbrace{\mathbb{E}^{\mathbb{P}} \left[\texttt{Payoff}\big(\texttt{Pol}, w(\cdot)\big) \right]}^{\text{mean payoff}}}_{\text{pessimistic over probabilities}}$$

Convex risk measures cover a wide range of risk criteria

- Different probabilities $\mathbb{P}$, termed as beliefs or priors and belonging to a set $\mathcal{P}$ of admissible probabilities on Ω
- To each probability $\mathbb{P}$ is attached a plausibility $\Theta(\mathbb{P})$

$$\max_{\texttt{Pol}\in\mathcal{U}^{ad}} \underbrace{\min_{\mathbb{P}\in\mathcal{P}} \overbrace{\mathbb{E}^{\mathbb{P}}\Big[\texttt{Payoff}\big(\texttt{Pol}, w(\cdot)\big)\Big]}^{\text{mean payoff}} - \overbrace{\Theta(\mathbb{P})}^{\text{plausibility}}}_{\text{pessimistic over probabilities}}$$

Non convex risk measures can lead to non diversification

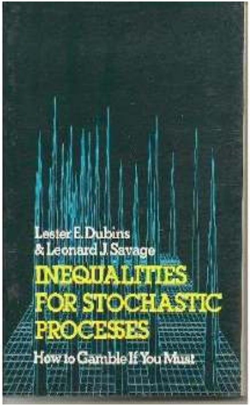

How to gamble if you must, L.E. Dubbins and L.J. Savage, 1965

Imagine yourself at a casino with \$1,000. For some reason, you desperately need \$10,000 by morning; anything less is worth nothing for your purpose.

The only thing possible is to gamble away your last cent, if need be, in an attempt to reach the target sum of \$10,000.

- The question is how to play, not whether. What ought you do? How should you play?
  - Diversify, by playing 1 \$ at a time?
  - Play boldly and concentrate, by playing 10,000 \$ only one time?
- What is your decision criterion?

Savage's minimal regret criterion... "Had I known"

$$\min_{\texttt{Pol}\in\mathfrak{U}^{ad}} \Big\{ \overbrace{\max_{w(\cdot)\in\Omega} \Big[\underbrace{\max_{\text{anticipative policies } \overline{\texttt{Pol}}} \texttt{Payoff}\big(\overline{\texttt{Pol}}, w(\cdot)\big) - \texttt{Payoff}\big(\texttt{Pol}, w(\cdot)\big)}_{\text{regret}} \Big]}^{\text{worst regret}} \Big\}$$

- If the DM knows the future in advance, she solves $\max_{\text{anticipative policies } \overline{\texttt{Pol}}} \texttt{Payoff}\big(\overline{\texttt{Pol}}, w(\cdot)\big)$, for each scenario $w(\cdot) \in \Omega$
- The regret attached to a non-anticipative policy $\texttt{Pol} \in \mathfrak{U}^{ad}$ is the loss due to not being visionary
- The best a non-visionary DM can do with respect to regret is minimizing it

Outline of the presentation

Upon which online information are decisions made?

We navigate between two stumbling blocks: rigidity and wizardry

- On the one hand, it is suboptimal to restrict oneself, as in the deterministic case, to open-loop controls depending only upon time, thereby ignoring the available information at the moment of making a decision
- On the other hand, it is impossible to suppose that we know in advance what will happen for all times: clairvoyance is impossible as well as look-ahead solutions

The in-between is **non-anticipativity constraint**

There are two ways to express the non-anticipativity constraint

Denote the uncertainties at time t by $w(t)$, and the control by $u(t)$

- ▷ Functional approach
  The control $u(t)$ may be looked after under the form

$$u(t) = \phi_t\big(\underbrace{w(t_0), \ldots, w(t-1)}_{\text{past}}\big)$$

  where ϕ_t is a function, called policy, strategy or decision rule

- ▷ Algebraic approach
  When uncertainties are considered as random variables (measurable mappings), the above formula for $u(t)$ expresses the measurability of the control variable $u(t)$ with respect to the past uncertainties, also written as

$$\sigma(u(t)) \subset \sigma\big(\underbrace{w(t_0), \ldots, w(t-1)}_{\text{past}}\big)$$

What is a solution at time t?

- In deterministic control, the solution $u(t)$ at time t is a single number
- In stochastic control, the solution $u(t)$ at time t is a random variable expressed
  - either as $u(t) = \phi_t\big(w(t_0), \ldots, w(t-1)\big)$, where $\phi_t : \mathbb{W}^{t-t_0} \to \mathbb{R}$
  - or as $u(t) : \Omega \to \mathbb{R}$ with measurability constraint $\sigma(u(t)) \subset \sigma\big(w(t_0), \ldots, w(t-1)\big)$
- Now, as time t goes on, the domain of the function ϕ_t expands, and so do the conditions $\sigma(u(t)) \subset \sigma\big(w(t_0), \ldots, w(t-1)\big)$
- Therefore, for numerical reasons, the information $\big(w(t_0), \ldots, w(t-1)\big)$ has to be compressed or approximated

Scenarios can be organized like a tree

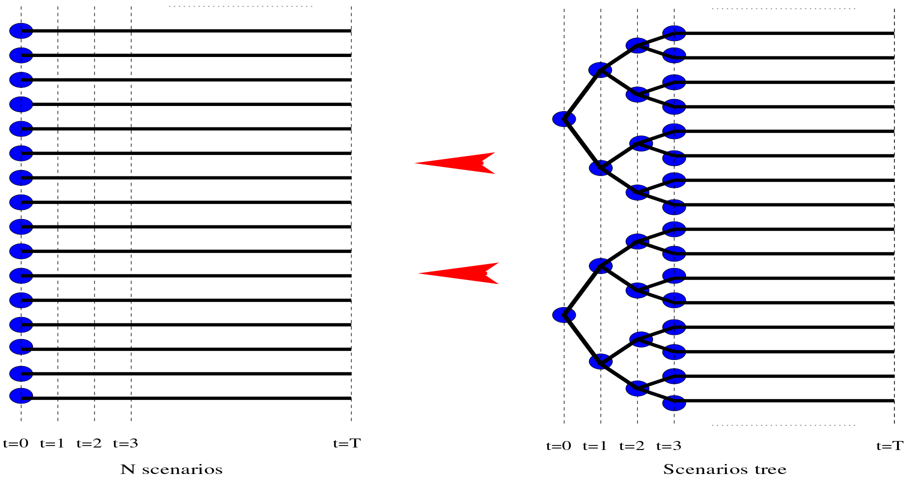

There are two classical ways to compress information

▷ State-based functional approach

In the special case of the Markovian framework with $\big(w(t_0), \dots, w(T)\big)$ white noise, there is no loss of optimality to look for solutions as

$$u(t) = \psi_t \underbrace{\big(x(t)\big)}_{\text{state}} \quad \text{where} \quad \underbrace{x(t) \in \mathbb{X}}_{\text{fixed space}}, \quad \underbrace{x(t+1) = F_t(x(t), u(t), w(t))}_{\text{dynamical equation}}$$

▷ Scenario-based measurability approach

- Scenarios are approximated by a finite family $\big(w^s(t_0), \dots, w^s(T)\big)$, $s \in S$
- Solutions $u^s(t)$ are indexed by $s \in S$ with the constraint that if two scenarios coincide up to time t, so must do the controls at time t

$$\big(w^s(t_0), \dots, w^{s'}(t-1)\big) = \big(w^{s'}(t_0), \dots, w^{s'}(t-1)\big) \Rightarrow u^s(t) = u^{s'}(t)$$

- In the case of the scenario tree approach, the scenarios $\big(w^s(t_0), \dots, w^s(T)\big)$, $s \in S$, are organized in a tree, and controls $u^n(t)$ are indexed by nodes n on the tree

More on what is a solution at time t
State-based approach $u(t) = \psi_t(x(t))$

- The mapping ψ_t can be computed in advance (that is, at initial time t_0) and evaluated at time t on the available online information at that time t
  - either exactly (for example, by dynamic programming)
  - or approximately (for example, among linear decision rules) because the computational burden of finding *any* function is heavy
- The value $u(t) = \psi_t(x(t))$ can be computed at time t
  - either exactly by solving a proper optimization problem, which raises issues of dynamic consistency
  - or approximately (for example, by assuming that controls from time t on are open-loop)

More on what is a solution at time t Scenario-based approach

- An optimal "solution" can be computed scenario by scenario, with the problem that we obtain solutions such that

$$\big(w^s(t_0), \dots, w^s(t-1)\big) = \big(w^{s'}(t_0), \dots, w^{s'}(t-1)\big) \text{ and } u^s(t) \neq u^{s'}(t)$$

- Optimal solutions can be computed scenario by scenario and then merged (for example, by Progressive Hedging) to be forced to satisfy

$$\big(w^s(t_0), \dots, w^s(t-1)\big) = \big(w^{s'}(t_0), \dots, w^{s'}(t-1)\big) \Rightarrow u^s(t) = u^{s'}(t)$$

- The value $u(t)$ can be computed at time t depending on $\big(w^s(t_0), \dots, w^s(t-1)\big)$
  - either exactly by solving a proper optimization problem, which raises issues of dynamic consistency
  - or approximately (for example, by a sequence of two-stages problems)

Outline of the presentation

Where do we stand?

- ▷ How one frames the non-anticipativity constraint impacts numerical resolution methods
- ▷ On a finite scenario space, one obtains large (deterministic) optimization problems on a tree
- ▷ Or large (deterministic) optimization problems indexed by scenarios
- ▷ Else, you resort to state-based formulations, with solutions as policies (dynamic programming)

Optimization approaches to attack complexity

Linear programming

- linear equations and inequalities
- no curse of dimension

Stochastic programming

- no special treatment of time and uncertainties
- no independence assumption
- decisions are indexed by a scenario tree
- what if information is not a node in the tree?

State-based dynamic optimization

- nonlinear equations and inequalities
- curse of dimensionality
- independence assumption on uncertainties
- special treatment of time (dynamic programming equation)
- decisions are indexed by an information state (feedback synthesis)
- an information state summarizes past controls and uncertainties
- decomposition-coordination methods to overcome the curse of dimensionality?

Summary

- *Stochastic* optimization highlights risk attitudes tackling
- Stochastic *dynamic* optimization emphasizes the handling of online information
- Many issues are raised, because
  - many ways to represent risk (criterion, constraints)
  - many information structures
  - tremendous numerical obstacles to overcome
- Each method has its numerical wall
  - in dynamic programming, the bottleneck is the dimension of the state (no more than 3)
  - in stochastic programming, the bottleneck is the number of stages (no more than 2)

Outline of the presentation

Outline of the presentation

Decomposition-coordination: divide and conquer

- Spatial decomposition
  - multiple players with their local information
  - local / regional / national /supranational
- Temporal decomposition
  - A state is an information summary
  - Time coordination realized through Dynamic Programming, by value functions
  - Hard nonanticipativity constraints
- Scenario decomposition
  - Along each scenario, sub-problems are deterministic (powerful algorithms)
  - Scenario coordination realized through Progressive Hedging, by updating nonanticipativity multipliers
  - Soft nonanticipativity constraints

Coupling constraints: an overview

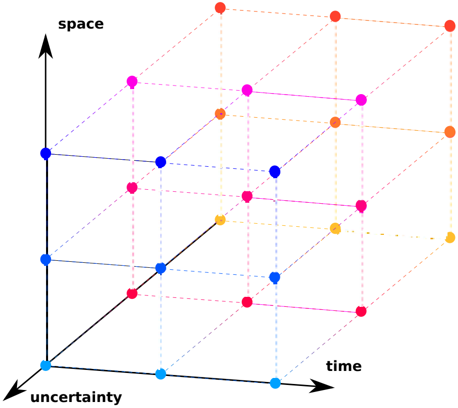

$$\min_{\mathbf{x},\,\mathbf{u}} \sum_{s=1}^{S} \sum_{i=1}^{N} \sum_{t=0}^{T} \pi_s L_{i,t}\big(\mathbf{x}_{i,t},\, \mathbf{u}_{i,t},\, \mathbf{w}_t\big)$$

Coupling constraints: time coupling

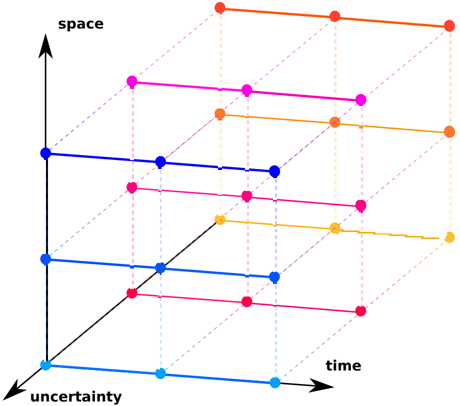

$$\min_{\mathbf{x},\,\mathbf{u}} \sum_{s=1}^{S} \sum_{i=1}^{N} \sum_{t=0}^{T} L_{i,t}\big(\mathbf{x}_{i,t},\, \mathbf{u}_{i,t},\, \mathbf{w}_t\big)$$

$$\text{s.t.} \quad \mathbf{x}_{i,t+1} = f_{i,t}(\mathbf{x}_{i,t}, \mathbf{u}_{i,t}, \mathbf{w}_t)$$

Coupling constraints: scenario coupling

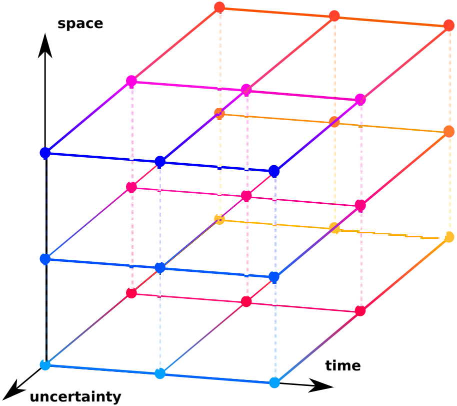

$$\min_{\mathbf{x},\,\mathbf{u}} \sum_{s=1}^{S} \sum_{i=1}^{N} \sum_{t=0}^{T} L_{i,t}\big(\mathbf{x}_{i,t},\, \mathbf{u}_{i,t},\, \mathbf{w}_t\big)$$

$$\text{s.t.} \quad \mathbf{x}_{i,t+1} = f_{i,t}(\mathbf{x}_{i,t}, \mathbf{u}_{i,t}, \mathbf{w}_t)$$

$$\sigma(\mathbf{u}_{i,t}) \subset \sigma(\mathbf{w}_0, \ldots, \mathbf{w}_t)$$

Coupling constraints: space coupling

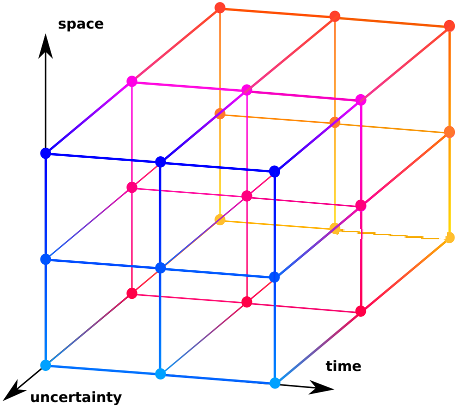

$$\min_{\mathbf{x},\,\mathbf{u}} \sum_{s=1}^{S}\sum_{i=1}^{N}\sum_{t=0}^{T} L_{i,t}\big(\mathbf{x}_{i,t},\, \mathbf{u}_{i,t},\, \mathbf{w}_t\big)$$

$$\text{s.t.} \quad \mathbf{x}_{i,t+1} = f_{i,t}(\mathbf{x}_{i,t}, \mathbf{u}_{i,t}, \mathbf{w}_t)$$

$$\sigma(\mathbf{u}_{i,t}) \subset \sigma(\mathbf{w}_0, \ldots, \mathbf{w}_t)$$

$$\sum_{i=1}^{N} \theta_{i,t}\big(\mathbf{x}_{i,t},\, \mathbf{u}_{i,t},\, \mathbf{w}_t\big) = 0$$

Decomposition/coordination methods: an overview

Main idea

1. decompose a large scale problem into smaller subproblems we are able to solve by efficient algorithms
2. coordinate the subproblems for the concatenation of their solutions to form the initial problem solution

How to decompose the problem by duality?

1. identify the coupling dimensions of the problem: time, uncertainty, space
2. dualize the coupling constraints by introducing multiplyers
3. split the problem into the resulting subproblems and coordinate them by means of the multiplyer

In the case of time decomposition, we can use the time arrow to chain static subproblems by the dynamics equation (without dualizing)

Decomposition/coordination methods: an overview

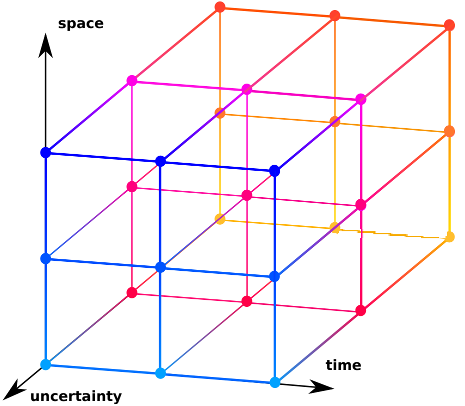

$$
\begin{aligned}
\min_{\mathbf{x},\,\mathbf{u}} \ & \mathbb{E}\Big(\sum_{i=1}^{N} \sum_{t=0}^{T} L_{i,t}(\mathbf{x}_{i,t},\, \mathbf{u}_{i,t},\, \mathbf{w}_t) \Big) \\
\text{s.t.} \ & \mathbf{x}_{i,t+1} = f_{i,t}(\mathbf{x}_{i,t}, \mathbf{u}_{i,t}, \mathbf{w}_t) \\
& \mathbf{u}_{i,t} = \mathbb{E}(\mathbf{u}_{i,t} \mid \mathbf{w}_0, \ldots, \mathbf{w}_t) \\
& \sum_{i=1}^{N} \theta_{i,t}(\mathbf{x}_{i,t},\, \mathbf{u}_{i,t},\, \mathbf{w}_t) = 0
\end{aligned}
$$

Decomposition/coordination methods: time coupling

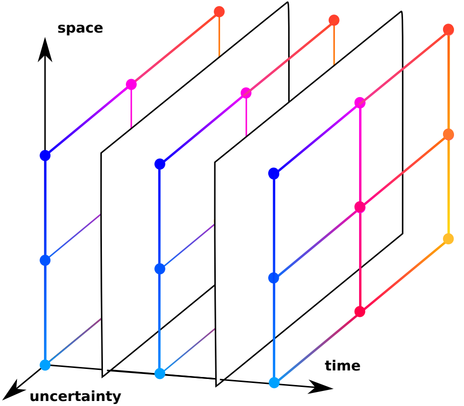

$$\min_{\mathbf{x},\,\mathbf{u}} \; \mathbb{E}\Big(\sum_{i=1}^{N} \sum_{t=0}^{T} L_{i,t}(\mathbf{x}_{i,t},\, \mathbf{u}_{i,t},\, \mathbf{w}_t) \Big)$$

s.t. $\mathbf{x}_{i,t+1} = f_{i,t}(\mathbf{x}_{i,t}, \mathbf{u}_{i,t}, \mathbf{w}_t)$

$$\mathbf{u}_{i,t} = \mathbb{E}(\mathbf{u}_{i,t} \mid \mathbf{w}_0, \ldots, \mathbf{w}_t)$$

$$\sum_{i=1}^{N} \theta_{i,t}(\mathbf{x}_{i,t},\, \mathbf{u}_{i,t},\, \mathbf{w}_t) = 0$$

[Stochastic Pontryagin]
[Dynamic Programming]

Decomposition/coordination methods: scenario coupling

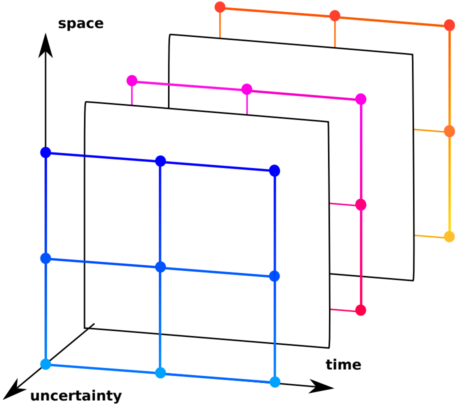

$$\min_{\mathbf{x},\,\mathbf{u}} \ \mathbb{E}\Bigg(\sum_{i=1}^{N}\sum_{t=0}^{T} L_{i,t}\big(\mathbf{x}_{i,t},\ \mathbf{u}_{i,t},\ \mathbf{w}_t\big)\Bigg)$$

$$\text{s.t.}\quad \mathbf{x}_{i,t+1} = f_{i,t}(\mathbf{x}_{i,t}, \mathbf{u}_{i,t}, \mathbf{w}_t)$$

$$\mathbf{u}_{i,t} = \mathbb{E}\big(\mathbf{u}_{i,t} \mid \mathbf{w}_0, \dots, \mathbf{w}_t\big)$$

$$\sum_{i=1}^{N} \theta_{i,t}\big(\mathbf{x}_{i,t},\ \mathbf{u}_{i,t},\ \mathbf{w}_t\big) = 0$$

[Progressive Hedging]

Rockafellar, R.T., Wets R. J-B.
Scenario and policy aggregation in optimization under uncertainty,
Mathematics of Operations Research, 16, pp. 119-147, 1991

Decomposition/coordination methods: space coupling

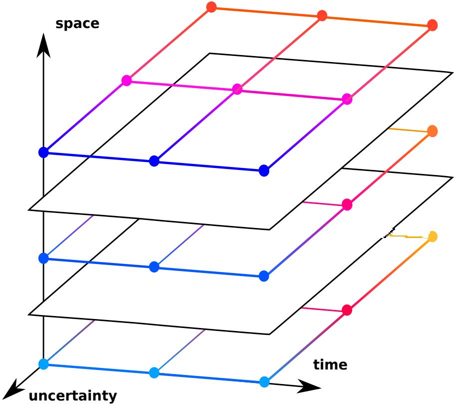

$$\min_{\mathbf{x},\,\mathbf{u}} \ \mathbb{E}\left(\sum_{i=1}^{N} \sum_{t=0}^{T} L_{i,t}\big(\mathbf{x}_{i,t},\, \mathbf{u}_{i,t},\, \mathbf{w}_t\big) \right)$$

$$\text{s.t.} \quad \mathbf{x}_{i,t+1} = f_{i,t}(\mathbf{x}_{i,t}, \mathbf{u}_{i,t}, \mathbf{w}_t)$$

$$\sigma(\mathbf{u}_{i,t}) \subset \sigma(\mathbf{w}_0, \ldots, \mathbf{w}_t)$$

$$\sum_{i=1}^{N} \theta_{i,t}(\mathbf{x}_{i,t},\, \mathbf{u}_{i,t},\, \mathbf{w}_t) = 0$$

[Our purpose now]

We have a nice decomposed problem but...

Flower structure

We are almost in the case where units could be driven independently one from another

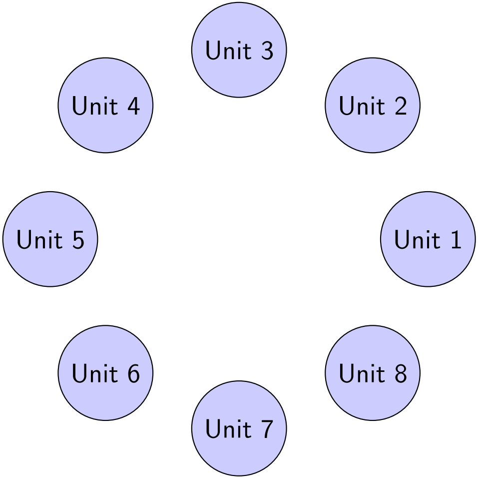

We have a nice decomposed problem but. . .

Flower structure

Unfortunately...

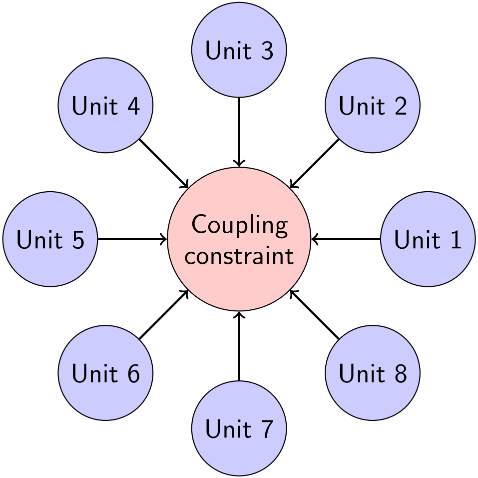

The associated optimization problem can be written as

$$\underbrace{\min_{(u_1,\ldots,u_N)} \sum_{i=1}^{N} J_i(u_i)}_{\text{costs minimization}} \quad \text{under} \quad \underbrace{\sum_{i=1}^{N} \Theta_i(u_i) = D}_{\text{supply = demand}}$$

where

- u_i is the decision of each unit i
- $J_i(u_i)$ is the cost of making decision u_i for unit i
- $\Theta_i(u_i)$ is the production induced by making decision u_i for unit i

Under appropriate duality assumptions, the associated optimization problem can be written without constraints

- For a proper Lagrange multiplier λ

$$\min_{(u_1,\ldots,u_N)} \sum_{i=1}^{N} J_i(u_i) + \lambda \underbrace{\Big(\sum_{i=1}^{N} \Theta_i(u_i) - D \Big)}_{\text{constraint}}$$

- We distribute the coupling constraint to each unit i

$$\min_{(u_1,\ldots,u_N)} \Big(\sum_{i=1}^{N} J_i(u_i) + \lambda \Theta_i(u_i) \Big) - \lambda D$$

- The problems splits into N optimization problems

$$\min_{u_i} \big(J_i(u_i) + \lambda \Theta_i(u_i) \big) , \quad \forall i = 1, \ldots, N$$

Proper prices allow decentralization of the optimum

$$\min_{(u_1,\ldots,u_N)} \sum_{i=1}^{N} J_i(u_i) \quad \text{under} \quad \sum_{i=1}^{N} \Theta_i(u_i) = D$$

The simplest decomposition/coordination scheme consists in

- ▷ buying the production of each unit at a price $\lambda^{(k)}$ at iteration k
- ▷ and letting each unit minimize its modified costs

$$\min_{u_i} J_i(u_i) + \underbrace{\lambda^{(k)}}_{\text{price}} \Theta_i(u_i)$$

- ▷ then, updating the price depending on the coupling constraint

$$\lambda^{(k+1)} = \lambda^{(k)} + \rho\Big(\sum_{i=1}^{N} \Theta_i(u_i) - D\Big)$$

(like in the "tâtonnement de Walras" in Economics)

What are the stakes if we extend spatial coupling constraint decomposition to the dynamical and stochastic setting?

- Allowing for time and uncertainties, we classically consider the criterion

$$\min_{\{\mathbf{u}_{i,t}\}_{i\in\{1,N\}}^{t\in\{0,T-1\}}} \mathbb{E}\left(\sum_{i=1}^{N}\left(\sum_{t=0}^{T-1} L_{i,t}\big(\mathbf{x}_{i,t},\mathbf{u}_{i,t},\mathbf{w}_{i,t}\big) + K_i\big(\mathbf{x}_{i,T}\big)\right)\right)$$

- under the constraints

$$\sum_{i=1}^{N}\Theta_{i,t}\big(\mathbf{x}_{i,t},\mathbf{u}_{i,t},\mathbf{w}_{i,t}\big) - \mathbf{d}_t = 0, \qquad t\in[\![0, T-1]\!]$$

$$\mathbf{x}_{i,t+1} = f_{i,t}\big(\mathbf{x}_{i,t},\mathbf{u}_{i,t},\mathbf{w}_{i,t}\big), \qquad i\in[\![1,N]\!],\ t\in[\![0,T-1]\!]$$

To avoid finding "magical solutions", only implementable by a wizard knowing the future, we need to specify information constraints

- The optimization problem is not well posed, because we have not specified upon what depends the control $\mathbf{u}_{i,t}$ of each unit i at each time t
- In the causal and perfect memory case, we express that the control $\mathbf{u}_{i,t}$ depends of all past noises up to time t
  - either by a functional approach

$$\mathbf{u}_{i,t} = \phi_{i,t}\big(\mathbf{w}_{1,0}, \ldots, \mathbf{w}_{N,0}, \mathbf{d}_0 \ldots \ldots \mathbf{w}_{1,t}, \ldots, \mathbf{w}_{N,t}, \mathbf{d}_t\big)$$

  - or by an algebraic approach

$$\sigma(\mathbf{u}_{i,t}) \subset \sigma\big(\mathbf{w}_{1,0}, \ldots, \mathbf{w}_{N,0}, \mathbf{d}_0 \ldots \ldots \mathbf{w}_{1,t}, \ldots, \mathbf{w}_{N,t}, \mathbf{d}_t\big)$$

Looking after decentralizing prices models

- Going on with the previous scheme, each unit i solves

$$\min_{\mathbf{u}_{i,0},\dots,\mathbf{u}_{i,T-1}} \mathbb{E}\left(\sum_{t=0}^{T-1} \left(L_{i,t}\left(\mathbf{x}_{i,t}, \mathbf{u}_{i,t}, \mathbf{w}_{i,t}\right) + \underbrace{\lambda_{i,t}^{(k)}}_{\text{price}} \Theta_{i,t}\left(\mathbf{x}_{i,t}, \mathbf{u}_{i,t}, \mathbf{w}_{i,t}\right)\right) + K_i\left(\mathbf{x}_{i,T}\right)\right)$$

$$\mathbf{x}_{i,t+1} = f_{i,t}\left(\mathbf{x}_{i,t}, \mathbf{u}_{i,t}, \mathbf{w}_{i,t}\right), \quad t \in [\![0, T-1]\!]$$

- The optimal controls $\mathbf{u}_{i,t}^{\star}$ of this problem depend
  - upon the local state $\mathbf{x}_{i,t}$
  - and ... upon all past prices $(\lambda_{i,0}^{(k)}, \dots, \lambda_{i,t}^{(k)})$!
- Research axis: find an approximate dynamical model for the price process, driven by proper information; for instance, replace $\lambda_{i,t}^{(k)}$ by $\mathbb{E}\left(\lambda_{i,t}^{(k)} \mid \mathbf{y}_t\right)$, where the information variable $\mathbf{y}_t$ is a Markov process (short time memory) $\rightarrow$ "demand response", "adaptive tariffs",

Dual Approximate Dynamic Programming

Samples/scenarios of dual variable at iteration k

λ^k

Dual Approximate Dynamic Programming

Samples/scenarios of
dual variable
at iteration k

We solve subproblems
using $\mathbb{E}(\lambda^k | \mathbf{y})$
by Dynamic Programming

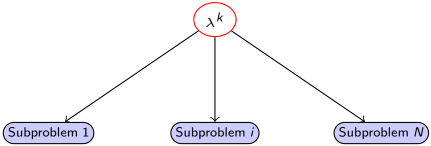

Dual Approximate Dynamic Programming

Samples/scenarios of dual variable at iteration k

We solve subproblems using $\mathbb{E}(\lambda^k | \mathbf{y})$ by Dynamic Programming

We obtain policies

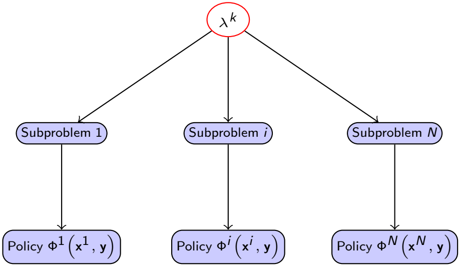

Dual Approximate Dynamic Programming

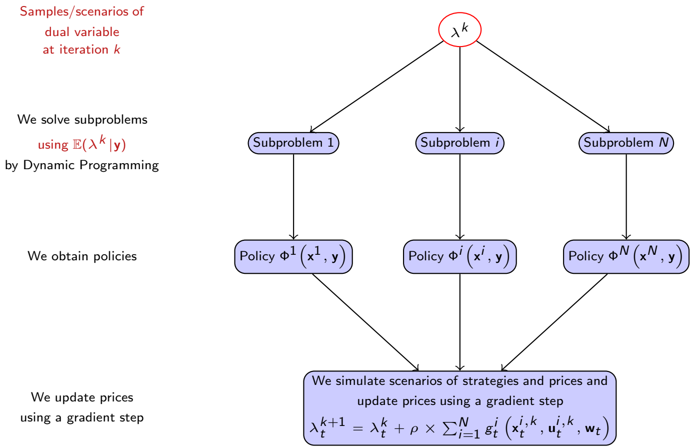

Dual Approximate Dynamic Programming

At iteration $k+1$

We solve subproblems
using $\mathbb{E}(\lambda^k|\mathbf{y})$
by Dynamic Programming

We obtain policies

We update prices
using a gradient step

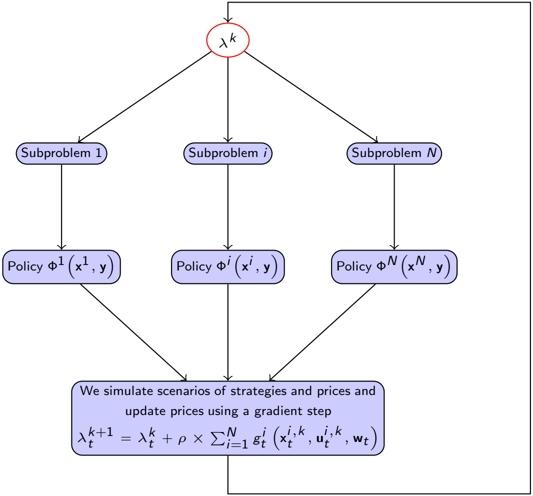

Extension to interconnected dams

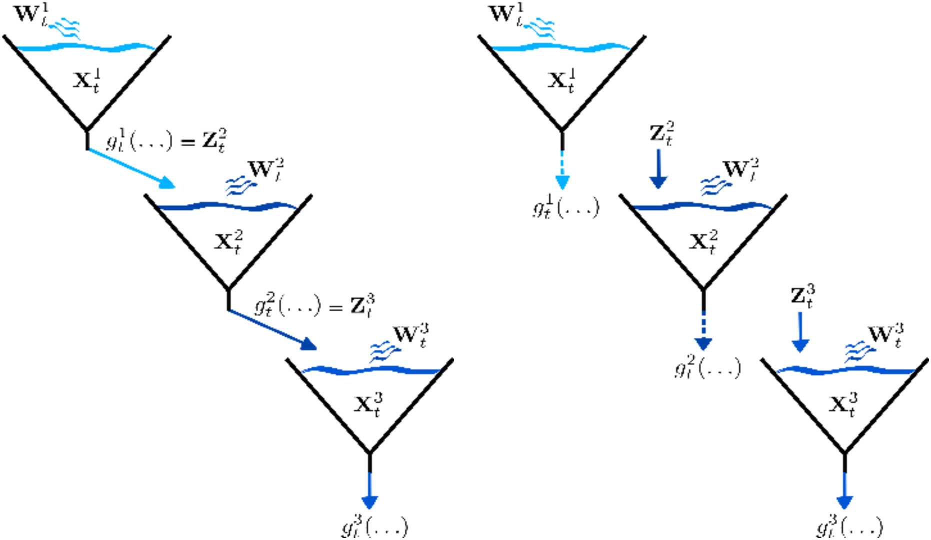

Contribution to dynamic tariffs

- ▷ Spatial decomposition of a dynamic stochastic optimization problem
- ▷ Lagrange multipliers attached to spatial coupling constraints are stochastic processes (prices)
- ▷ By projecting these prices, one expects to identify approximate dynamic models
- ▷ Such prices dynamic models are interpreted as dynamic tariffs

Outline of the presentation

1. Long term industry-academy cooperation
   - École des Ponts ParisTech–Cermics–Optimization and Systems
   - Industry partners of the Optimization and Systems Group
2. The remolding of power systems seen from an optimizer perspective
   - The remolding of power systems
   - Optimization is challenged
3. Moving from deterministic to stochastic dynamic optimization
   - Working out a toy example
   - Expliciting risk attitudes
   - Handling online information
   - Discussing framing and resolution methods
4. Two snapshots on ongoing research
   - Decomposition-coordination optimization methods under uncertainty
   - Risk constraints in optimization
5. A need for training and research

Tourism issues impose constraints upon traditional economic management of a hydro-electric dam

- ▷ Maximizing the revenue from turbinated water
- ▷ under a tourism constraint of having enough water in July and August

We consider a single dam nonlinear dynamical model in the decision-hazard setting

We model the dynamics of the water volume in a dam by

$$\underbrace{S(t+1)}_{\text{future volume}} = \min\{S^\sharp, \underbrace{S(t)}_{\text{volume}} - \underbrace{q(t)}_{\text{turbined}} + \underbrace{a(t)}_{\text{inflow volume}} \}$$

- $S(t)$ volume (stock) of water at the beginning of period $[t, t+1[$
- $q(t)$ turbined outflow volume during $[t, t+1[$
  - decided at the beginning of period $[t, t+1[$
  - chosen such that $0 \leq q(t) \leq \min\{S(t), q^\sharp\}$
- $a(t)$ inflow water volume (rain, etc.) during $[t, t+1[$, which materializes at the end $t+1$ of period $[t, t+1[$
- $S^\sharp$ dam capacity

The setting is called decision-hazard because the decision $q(t)$ is made before the hazard $a(t)$

The red stock trajectories fail to meet the tourism constraint in July and August

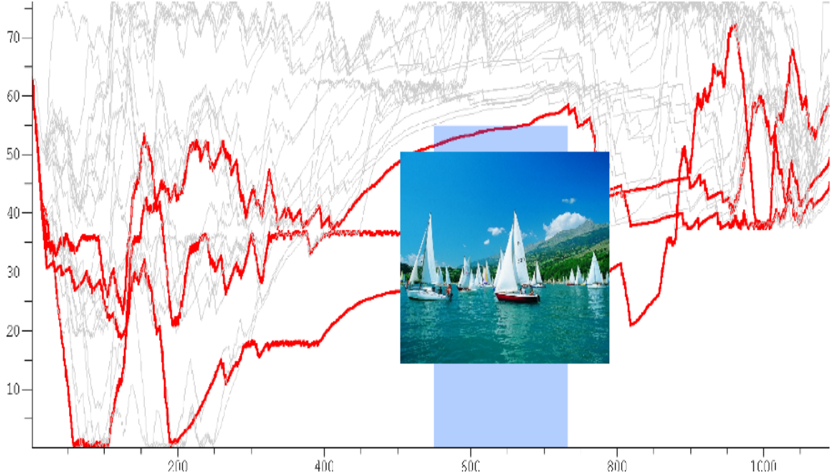

In the risk-neutral economic approach, an optimal management maximizes the expected payoff

- Suppose that
  - turbined water $q(t)$ is sold at price $p(t)$,
    related to the price at which energy can be sold at time t
  - a probability $\mathbb{P}$ is given on the set $\Omega = \mathbb{R}^{T-t_0} \times \mathbb{R}^{T-t_0}$
    of water inflows scenarios $\big(a(t_0), \ldots, a(T-1)\big)$
    and prices scenarios $\big(p(t_0), \ldots, p(T-1)\big)$
  - at the horizon, the final volume $S(T)$ has a value $\mathtt{K}\big(S(T)\big)$,
    the "final value of water"
- The traditional (risk-neutral) economic problem is to maximize the intertemporal payoff (without discounting if the horizon is short)

$$\max \mathbb{E}\left[\sum_{t=t_0}^{T-1}\left(\overbrace{p(t)}^{\text{price}}\ \overbrace{q(t)}^{\text{turbined}}\ \underbrace{-\epsilon q(t)^2}_{\text{turbined costs}}\right) + \overbrace{\mathtt{K}\big(S(T)\big)}^{\text{final volume utility}}\right]$$

We now have a stochastic optimization problem, where the tourism constraint still needs to be dressed in formal clothes

- Traditional cost minimization/payoff maximization

$$\max \mathbb{E}\left[\sum_{t=t_0}^{T-1} \overbrace{p(t)q(t) - \epsilon q(t)^2}^{\text{turbined water payoff}} + \overbrace{\mathtt{K}(S(T))}^{\text{final volume utility}}\right]$$

- Tourism constraint

$$\text{volume } S(t) \geq S^\flat \,, \; \forall t \in \mathcal{T} = \{\text{July, August}\}$$

- In what sense should we consider this inequality which involves the random variables $S(t)$ for $t \in \mathcal{T}$?

Robust / almost sure / probability constraint

- Robust constraints: for all the scenarios in a subset $\overline{\Omega} \subset \Omega$

$$S(t) \geq S^\flat \ , \ \ \forall t \in \mathcal{T}$$

- Almost sure constraints

$$\text{Probability} \left\{ S(t) \geq S^\flat \ , \ \ \forall t \in \mathcal{T} \right\} = 1$$

- Probability constraints, with "confidence" level $p \in [0, 1]$

$$\text{Probability} \left\{ S(t) \geq S^\flat \ , \ \ \forall t \in \mathcal{T} \right\} \geq p$$

- and also by penalization, or in the mean, etc.

Our problem may be clothed as a stochastic optimization problem under a probability constraint

$$P(T) = \sum_{t=t_0}^{T-1} \overbrace{p(t)q(t) - \epsilon q(t)^2}^{\text{turbined water payoff}} + \overbrace{K(S(T))}^{\text{final volume utility}}$$

- The traditional economic problem is $\max \mathbb{E}[P(T)]$
- and a failure tolerance is accepted

$$\text{Probability}\left\{S(t) \geq S^\flat\ ,\ \forall t \in \mathcal{T}\right\} \geq 90\%$$

- Details concerning the theoretical and numerical resolution are available on demand ;-)

Details concerning the theoretical and numerical resolution are available on demand ;-)

- $\pi(0) = 1$ and $\pi(t+1) = \begin{cases} \mathbf{1}_{\{S(t+1) \geq S^\flat\}} \times \pi(t) & \text{if } t \in \mathcal{T} \\ \pi(t) & \text{else} \end{cases}$
- $\mathbb{P}\left[S(\tau) \geq S^\flat \;,\; \forall \tau \in \mathcal{T}\right]$
  $= \mathbb{E}\left[\mathbf{1}_{\{S(\tau) \geq S^\flat \;,\; \forall \tau \in \mathcal{T}\}}\right]$
  $= \mathbb{E}\left[\prod_{\tau \in \mathcal{T}} \mathbf{1}_{\{S(\tau) \geq S^\flat\}}\right]$
  $= \mathbb{E}\left[\pi(T)\right]$

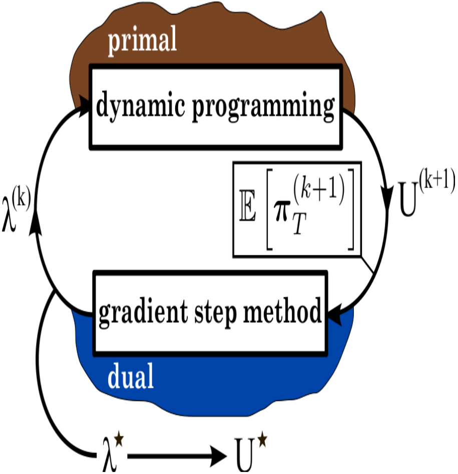

90% of the stock trajectories meet the tourism constraint

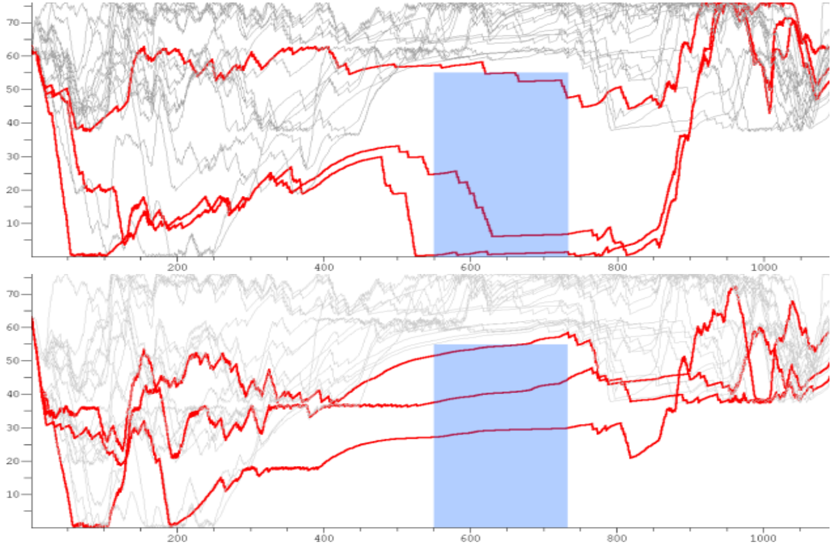

Our resolution approach brings a sensible improvement compared to standard procedures

OPTIMAL POLICIES	OPTIMIZATION		SIMULATION		
	Iterations	Time	Gain	Respect	Well behaviour
Standard	15	10 mn	**ref**	**0,9**	**no**
Convenient	**10**	**160 mn**	**-3.20%**	**0,9**	**yes**
Heuristic	10	160 mn	**-3.25%**	**0,9**	**yes**

However, though the expected payoff is optimal, the payoff effectively realized can be far from it

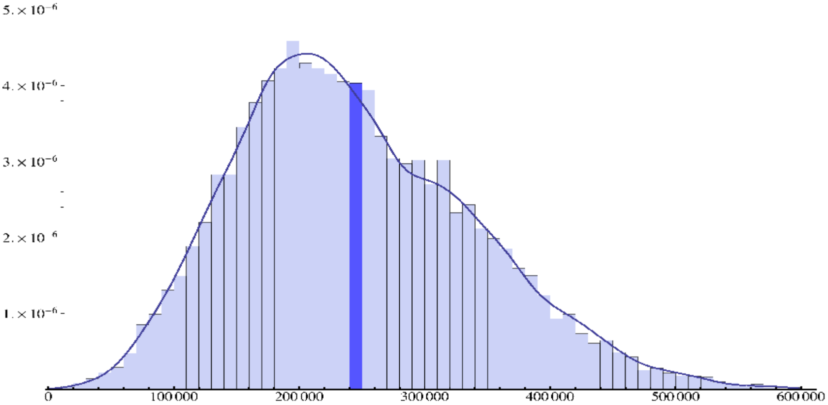

We propose a stochastic viability formulation to treat symmetrically and to guarantee both environmental and economic objectives

- Given two thresholds to be guaranteed
  - a volume $S^\flat$ (measured in cubic hectometers hm^3)
  - a payoff $P^\flat$ (measured in numeraire \$)
- we look after policies achieving the maximal viability probability

$$\Pi(S^\flat, P^\flat) = \max \text{Proba} \left\{ \begin{array}{l} \text{water inflow scenarios along which} \\ \text{the volumes } S(t) \geq S^\flat \\ \text{for all time } t \in \{ \text{July, August} \} \\ \text{and the final payoff } P(T) \geq P^\flat \end{array} \right\}$$

- $\Pi(S^\flat, P^\flat)$ is the maximal probability to guarantee to be above the thresholds $S^\flat$ and $P^\flat$

The stochastic viability formulation requires to redefine state and dynamics

- The state is the couple $x(t) = \big(S(t), P(t)\big)$ volume/payoff
- The control $u(t) = q(t)$ is the turbined water
- The dynamics is

$$\underbrace{S(t+1)}_{\text{future volume}} = \min\{S^{\sharp}, \underbrace{S(t)}_{\text{volume}} - \underbrace{q(t)}_{\text{turbined}} + \underbrace{a(t)}_{\text{inflow volume}} \} , \quad t = t_0, \ldots, T-1$$

$$\underbrace{P(t+1)}_{\text{future payoff}} = \underbrace{P(t)}_{\text{payoff}} + \underbrace{p(t)q(t) - \epsilon q(t)^2}_{\text{turbined water payoff}} , \quad t = t_0, \ldots, T-2$$

$$P(T) = P(T-1) + \underbrace{\mathtt{K}\big(S(T)\big)}_{\text{final volume utility}}$$

In the stochastic viability formulation, we dress objectives as state constraints

- ▷ The control constraints are

$$u(t) \in \mathbb{B}\big(t, x(t)\big) \iff 0 \leq q(t) \leq \min\{S(t), q^\sharp\}$$

- ▷ The state constraints are

$$x(t) \in \mathbb{A}(t) \iff \begin{cases} S(t) \geq S^\flat \\ P(T) \geq P^\flat \end{cases} , \quad \forall t \in \{\text{ July, August }\}$$

For each couple of thresholds on payoff and stock, we write a dynamic programming equation

- Abstract version

$$V(T,x) = \mathbf{1}_{\mathbb{A}(T)}(x)$$
$$V(t,x) = \mathbf{1}_{\mathbb{A}(t)}(x) \max_{u \in \mathbb{B}(t,x)} \mathbb{E}_{w(t)} \Big[V\Big(t+1, \mathtt{Dyn}\big(t,x,u,w(t)\big) \Big) \Big]$$

- Specific version

$$V(T,S,P) = \mathbf{1}_{\{P \geq P^{\flat}\}}$$
$$V(T-1,S,P) = \max_{0 \leq q \leq \min\{S,q^{\sharp}\}} \mathbb{E}_{a(T-1),p(T-1)} \Big[V\Big(t+1, S-q+a(t), P+\mathtt{K}(S) \Big) \Big]$$
$$V(t,S,P) = \max_{0 \leq q \leq \min\{S,q^{\sharp}\}} \mathbb{E}_{a(t),p(t)} \Big[V\Big(t+1, S-q+a(t), P+p(t)q-\epsilon q^2 \Big) \Big] ,$$
$$t \notin \{\text{ July, August }\}$$

$$V(t,S,P) = \mathbf{1}_{\{S \geq S^{\flat}\}} \max_{0 \leq q \leq \min\{S,q^{\sharp}\}} \mathbb{E}_{a(t),p(t)} \Big[V\Big(t+1, S-q+a(t), P+p(t)q-\epsilon q^2 \Big) \Big] ,$$
$$t \in \{\text{ July, August }\}$$

We plot iso-values for the maximal viability probability as a function of guaranteed thresholds $S^\flat$ and $P^\flat$

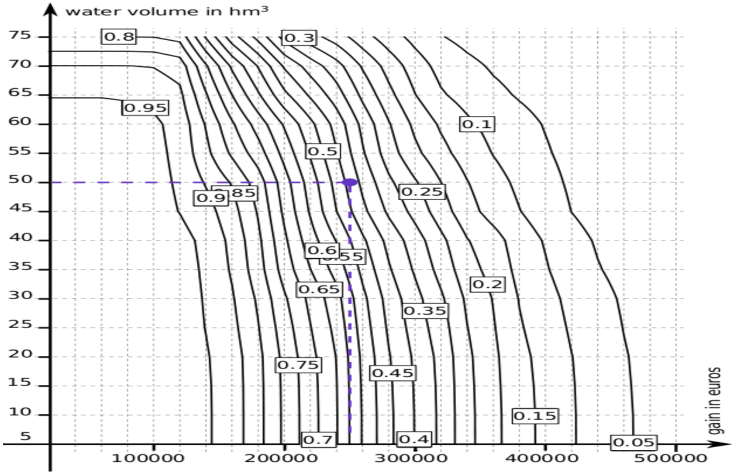

The probability distribution of the random gain reflects the viability objectives

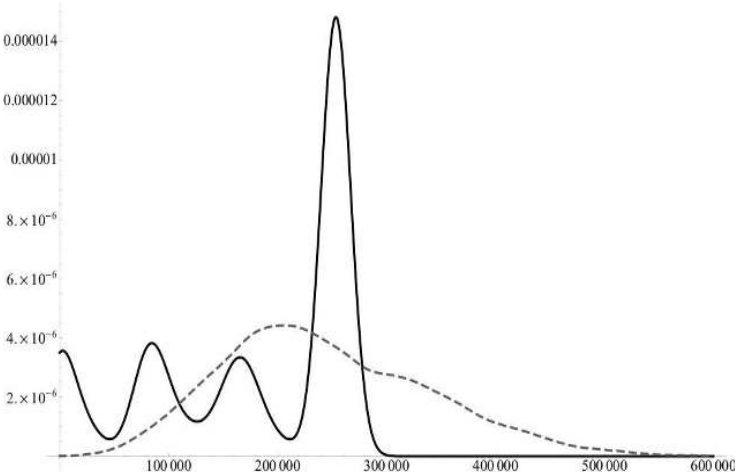

Contribution to quantitative sustainable management

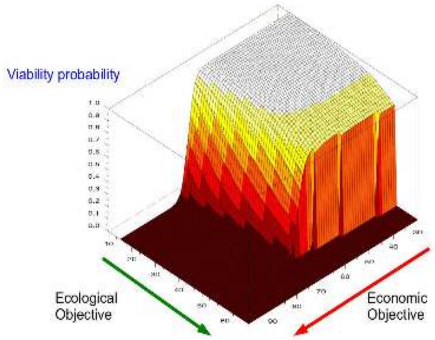

- ▷ Conceptual framework for quantitative sustainable management
- ▷ Managing ecological and economic conflicting objectives
- ▷ Displaying tradeoffs between ecology and economy sustainability thresholds and risk

Outline of the presentation

1. Long term industry-academy cooperation
2. The remolding of power systems seen from an optimizer perspective
3. Moving from deterministic to stochastic dynamic optimization
4. Two snapshots on ongoing research
5. A need for training and research

Trends are favorable to statistics and optimization

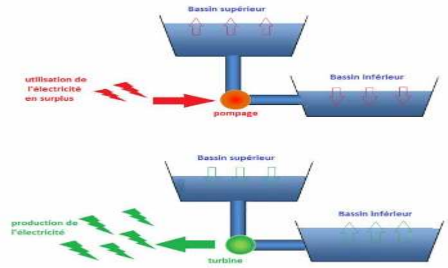

- More telecom technology
  ↪ more data
- More data, more unpredicability
  ↪ more statistics
- More unpredicability
  ↪ more storage
  ↪ more dynamic optimization
- More unpredicability
  ↪ more stochastic dynamic optimization

A context of increasing complexity

- ▷ Multiple energy resources: photovoltaic, solar heating, heatpumps, wind, hydraulic power, combined heat and power
- ▷ Spatially distributed energy resources (onshore and offshore windpower, solarfarms), producers, consumers
- ▷ Strongly variable production: wind, solar
- ▷ Intermittent demand: electrical vehicles
- ▷ Two-ways flows in the electrical network
- ▷ Environmental and risk constraints ($CO2$, nuclear risk, land use)

Challenges ahead for stochastic optimization

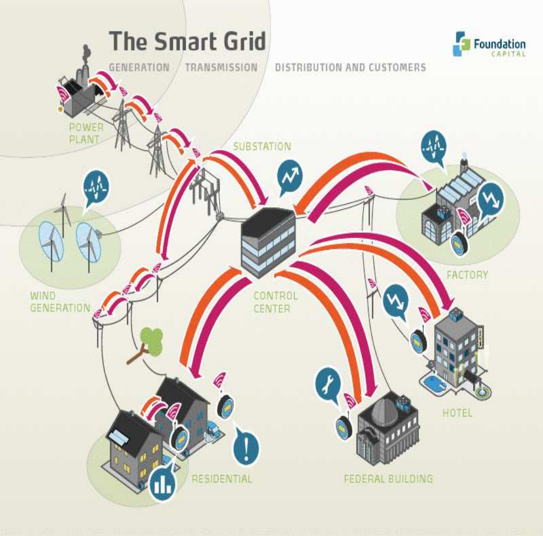

- ▷ large scale stochastic optimization
- ▷ various risk constraints
- ▷ decentralized and private information
- ▷ game theory, stochastic equilibrium, market design. . .

Spatial Decomposition/Coordination Methods for Stochastic Optimal Control Problems

Practical aspects and theoretical questions

J.-C. Alais, P. Carpentier, J-Ph. Chancelier,
M. De Lara, V. Leclère

École des Ponts ParisTech

3 November 2014

Large scale storage systems stand as powerful motivation

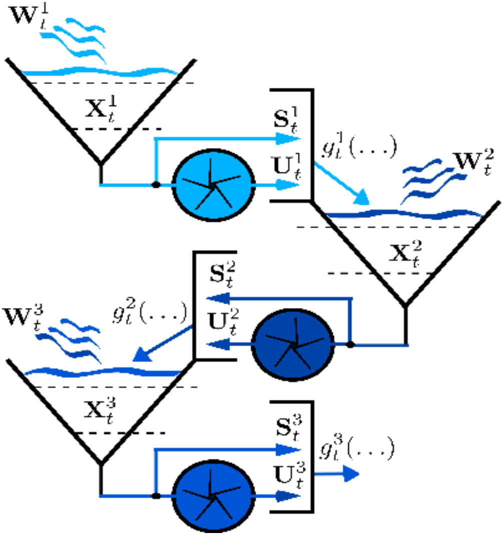

To make a long story short

We look after strategies as solutions of large scale stochastic optimal control problems,
for example, the optimal management over a given time horizon of a large amount of dynamical production units

- To obtain decision strategies (closed-loop controls), we use Dynamic Programming or related methods
  - **Assumption**: Markovian case
  - **Difficulty**: curse of dimensionality
- To use decomposition/coordination techniques, we have to deal with the information pattern of the stochastic optimization problem

A long-term effort in our group

1976 A. Benveniste, P. Bernhard, G. Cohen, "On the decomposition of stochastic control problems", *IRIA-Laboria research report*, No. 187, 1976.

1996 P. Carpentier, G. Cohen, J.-C. Culioli, A. Renaud, "Stochastic optimization of unit commitment: a new decomposition framework", *IEEE Transactions on Power Systems*, Vol. 11, No. 2, 1996.

2006 C. Strugarek, "Approches variationnelles et autres contributions en optimisation stochastique", *Thèse de l'ENPC*, mai 2006.

2010 K. Barty, P. Carpentier, P. Girardeau, "Decomposition of large-scale stochastic optimal control problems", *RAIRO Operations Research*, Vol. 44, No. 3, 2010.

2014 V. Leclère, "Contributions to decomposition methods in stochastic optimization", *Thèse de l'Université Paris-Est*, juin 2014.

Lecture outline

1. Decomposition and coordination
   - A bird's eye view of decomposition methods
   - (A brief insight into Progressive Hedging)
   - Spatial decomposition methods in the deterministic case
   - The stochastic case raises specific obstacles
2. Dual approximate dynamic programming (DADP)
   - Problem statement
   - DADP principle and implementation
   - Numerical results on a small size problem
3. Theoretical questions
   - Existence of a saddle point
   - Convergence of the Uzawa algorithm
   - Convergence w.r.t. information
4. Conclusion

Decomposition-coordination: divide and conquer

- Spatial decomposition
  - Multiple players with their local information
  - Scales: local / regional / national /supranational
- Temporal decomposition
  - A state is an information summary
  - Time coordination realized through Dynamic Programming, by value functions
  - Hard nonanticipativity constraints
- Scenario decomposition
  - Along each scenario, sub-problems are deterministic (powerful algorithms)
  - Scenario coordination realized through Progressive Hedging, by updating nonanticipativity multipliers
  - Soft nonanticipativity constraints

Couplings for stochastic problems

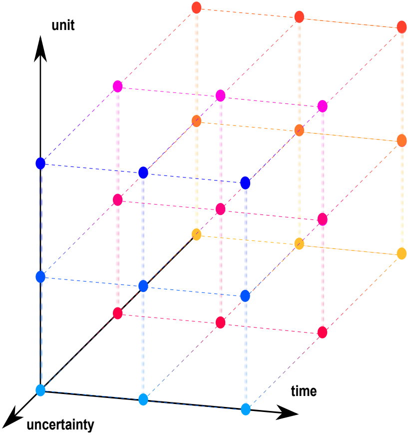

$$\min \sum_{\omega} \sum_{i} \sum_{t} \pi_\omega L_t^i(\mathbf{x}_t^i, \mathbf{u}_t^i, \mathbf{w}_{t+1})$$

Couplings for stochastic problems: in time

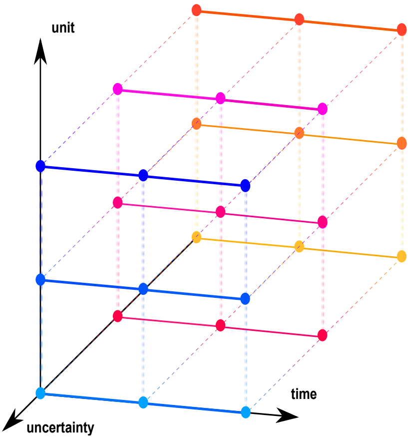

$$\min \sum_{\omega} \sum_{i} \sum_{t} \pi_\omega L_t^i(\mathbf{x}_t^i, \mathbf{u}_t^i, \mathbf{w}_{t+1})$$

s.t. $\mathbf{x}_{t+1}^i = f_t^i(\mathbf{x}_t^i, \mathbf{u}_t^i, \mathbf{w}_{t+1})$

Couplings for stochastic problems: in uncertainty

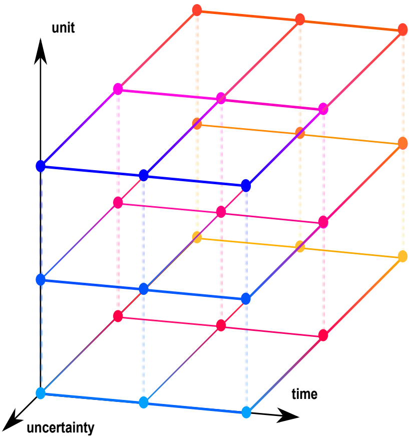

$$\min \sum_{\omega} \sum_{i} \sum_{t} \pi_\omega L_t^i(\mathbf{x}_t^i, \mathbf{u}_t^i, \mathbf{w}_{t+1})$$

$$\text{s.t.} \quad \mathbf{x}_{t+1}^i = f_t^i(\mathbf{x}_t^i, \mathbf{u}_t^i, \mathbf{w}_{t+1})$$

$$\mathbf{u}_t^i = \mathbb{E}\left(\mathbf{u}_t^i \;\middle|\; \mathbf{w}_1, \ldots, \mathbf{w}_t\right)$$

Couplings for stochastic problems: in space

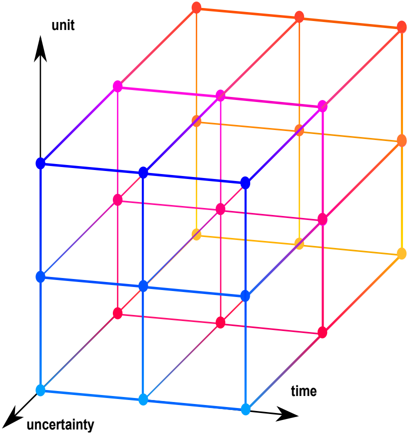

$$\min \sum_{\omega} \sum_{i} \sum_{t} \pi_{\omega} L_t^i(\mathbf{x}_t^i, \mathbf{u}_t^i, \mathbf{w}_{t+1})$$

$$\text{s.t.} \quad \mathbf{x}_{t+1}^i = f_t^i(\mathbf{x}_t^i, \mathbf{u}_t^i, \mathbf{w}_{t+1})$$

$$\mathbf{u}_t^i = \mathbb{E}\left(\mathbf{u}_t^i \;\middle|\; \mathbf{w}_1, \dots, \mathbf{w}_t\right)$$

$$\sum_{i} \theta_t^i(\mathbf{x}_t^i, \mathbf{u}_t^i) = 0$$

Can we decouple stochastic problems?

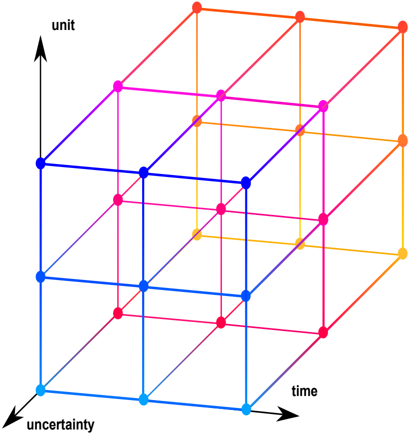

$$\min \sum_{\omega} \sum_{i} \sum_{t} \pi_\omega L_t^i(\mathbf{x}_t^i, \mathbf{u}_t^i, \mathbf{w}_{t+1})$$

$$\text{s.t.} \quad \mathbf{x}_{t+1}^i = f_t^i(\mathbf{x}_t^i, \mathbf{u}_t^i, \mathbf{w}_{t+1})$$

$$\mathbf{u}_t^i = \mathbb{E}\left(\mathbf{u}_t^i \;\middle|\; \mathbf{w}_1, \ldots, \mathbf{w}_t\right)$$

$$\sum_i \theta_t^i(\mathbf{x}_t^i, \mathbf{u}_t^i) = 0$$

Decompositions for stochastic problems: in time

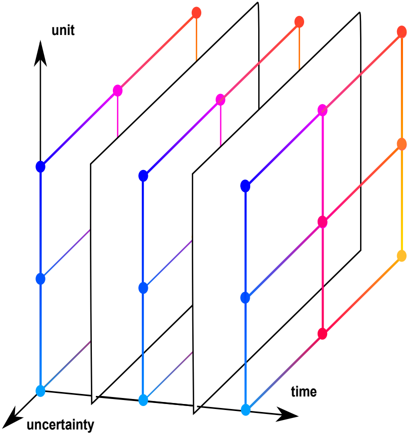

$$\min \sum_{\omega} \sum_{i} \sum_{t} \pi_{\omega} L_t^i(\mathbf{x}_t^i, \mathbf{u}_t^i, \mathbf{w}_{t+1})$$

s.t. $\mathbf{x}_{t+1}^i = f_t^i(\mathbf{x}_t^i, \mathbf{u}_t^i, \mathbf{w}_{t+1})$

$$\mathbf{u}_t^i = \mathbb{E}\left(\mathbf{u}_t^i \,\middle|\, \mathbf{w}_1, \ldots, \mathbf{w}_t\right)$$

$$\sum_{i} \theta_t^i(\mathbf{x}_t^i, \mathbf{u}_t^i) = 0$$

Dynamic Programming

Bellman (56)

Decompositions for stochastic problems: in uncertainty

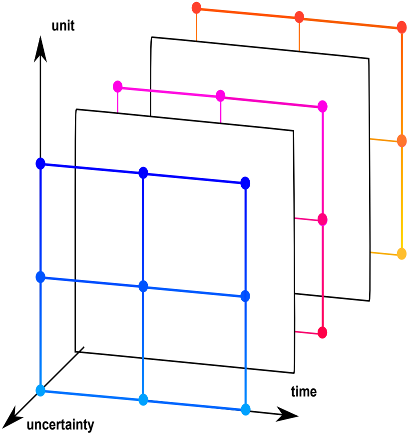

$$\min \sum_{\omega} \sum_{i} \sum_{t} \pi_{\omega} L_t^i(\mathbf{x}_t^i, \mathbf{u}_t^i, \mathbf{w}_{t+1})$$

s.t. $\mathbf{x}_{t+1}^i = f_t^i(\mathbf{x}_t^i, \mathbf{u}_t^i, \mathbf{w}_{t+1})$

$$\mathbf{u}_t^i = \mathbb{E}\left(\mathbf{u}_t^i \;\middle|\; \mathbf{w}_1, \dots, \mathbf{w}_t \right)$$

$$\sum_{i} \theta_t^i(\mathbf{x}_t^i, \mathbf{u}_t^i) = 0$$

Progressive Hedging

Rockafellar - Wets (91)

Decompositions for stochastic problems: in space

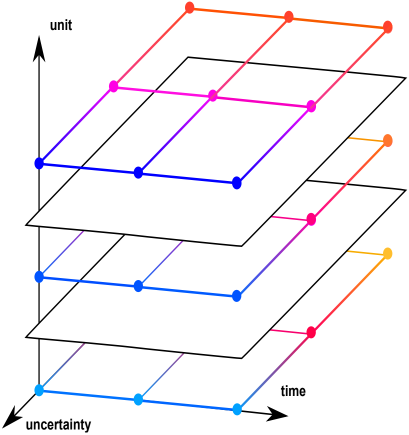

$$\min \sum_{\omega} \sum_{i} \sum_{t} \pi_\omega L_t^i(\mathbf{x}_t^i, \mathbf{u}_t^i, \mathbf{w}_{t+1})$$

$$\text{s.t.} \quad \mathbf{x}_{t+1}^i = f_t^i(\mathbf{x}_t^i, \mathbf{u}_t^i, \mathbf{w}_{t+1})$$

$$\mathbf{u}_t^i = \mathbb{E}\left(\mathbf{u}_t^i \;\middle|\; \mathbf{w}_1, \ldots, \mathbf{w}_t \right)$$

$$\sum_{i} \theta_t^i(\mathbf{x}_t^i, \mathbf{u}_t^i) = 0$$

Dual Approximate
Dynamic Programming

Outline of the presentation

1. Decomposition and coordination
   - A bird's eye view of decomposition methods
   - (A brief insight into Progressive Hedging)
   - Spatial decomposition methods in the deterministic case
   - The stochastic case raises specific obstacles
2. Dual approximate dynamic programming (DADP)
   - Problem statement
   - DADP principle and implementation
   - Numerical results on a small size problem
3. Theoretical questions
   - Existence of a saddle point
   - Convergence of the Uzawa algorithm
   - Convergence w.r.t. information
4. Conclusion

Non-anticipativity constraints are linear

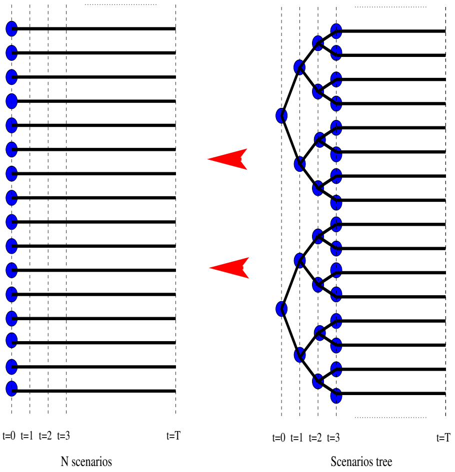

- From tree to scenarios (comb)
- Equivalent formulations of the non-anticipativity constraints
  - pairwise equalities
  - all equal to their mathematical expectation
- Linear structure

$$\mathbf{u}_t = \mathbb{E}\left(\mathbf{u}_t \,\middle|\, \mathbf{w}_1, \ldots, \mathbf{w}_t \right)$$

Progressive Hedging stands as a scenario decomposition method by dualizing the non-anticipativity constraints

- When the criterion is strongly convex,
  we use an algorithm "à la Uzawa"
  to obtain a scenario decomposition
- When the criterion is linear,
  Rockafellar - Wets (91) propose to use an augmented Lagrangian,
  and obtain the Progressive Hedging algorithm

Data: Initial multipliers $\big\{\{\lambda_t^{(0)}(\omega)\}_{t=0}^{T-1}\big\}_{\omega\in\Omega}$ and mean control $\big\{\overline{U}_n^{(0)}\big\}_{n\in\mathcal{T}}$;

Result: optimal feedback;

repeat

forall the *scenario* $\omega \in \Omega$ **do**

Solves the deterministic minimization problem for scenario ω with a measurability penalization, and obtain optimal control $\mathbf{u}^{(k+1)}$;

Update the mean controls

$$\overline{u}_n^{(k+1)} = \frac{\sum_{\omega\in n} \mathbf{u}_t^{(k+1)}(\omega)}{|n|}$$

Update the measurability penalization with

$$\lambda_t^{(k+1)}(\omega) = \lambda_t^{(k)}(\omega) + \rho\big(U_t(\omega)^{(k+1)} - \overline{u}_{n_t(\omega)}^{(k+1)}\big)$$

until $\mathbf{u}_t - \mathbb{E}\big(u_t^i \mid \mathbf{w}_1, \dots, \mathbf{w}_t\big) = 0$;

Decomposition and coordination

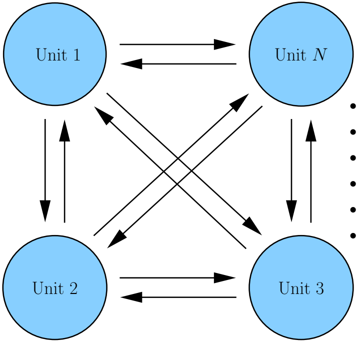

Interconnected units

- The system to be optimized consists of interconnected subsystems
- We want to use this structure to formulate optimization subproblems of reasonable complexity
- But the presence of interactions requires a level of coordination
- Coordination iteratively provides a local model of the interactions for each subproblem
- We expect to obtain the solution of the overall problem by concatenation of the solutions of the subproblems

Example: the "flower model"

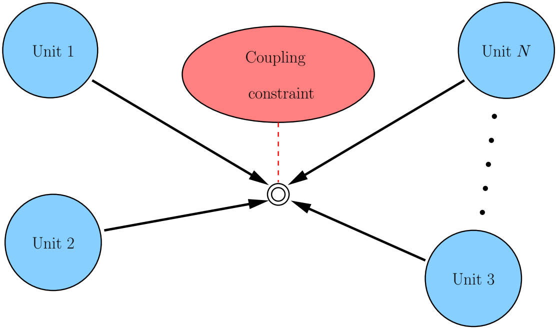

$$\min_u \sum_{i=1}^{N} J_i(u_i)$$

$$\text{s.t.} \quad \sum_{i=1}^{N} \theta_i(u_i) = \theta$$

Unit Commitment Problem

Intuition of spatial decomposition

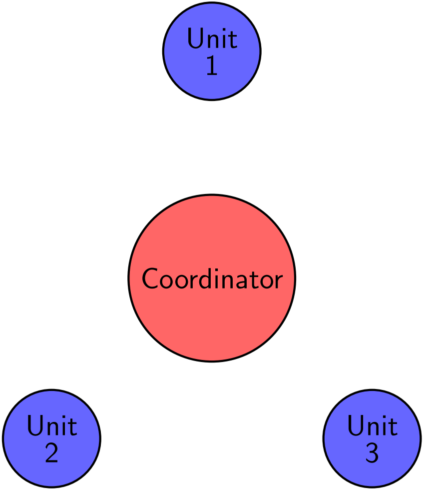

- Purpose: satisfy a demand with N production units, at minimal cost
- Price decomposition
  - the coordinator sets a price λ_t
  - the units send their production $\mathbf{u}_t^{(i)}$
  - the coordinator compares total production and demand, and then updates the price
  - and so on...

Intuition of spatial decomposition

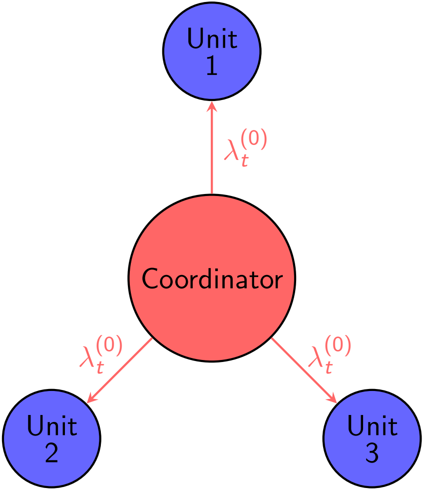

- Purpose: satisfy a demand with N production units, at minimal cost
- Price decomposition
  - the coordinator sets a price λ_t
  - the units send their production $\mathbf{u}_t^{(i)}$
  - the coordinator compares total production and demand, and then updates the price
  - and so on...

Intuition of spatial decomposition

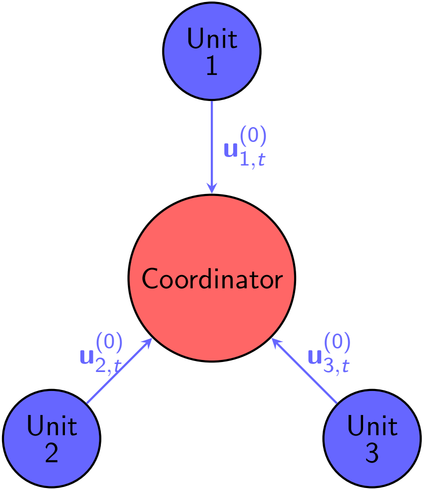

- Purpose: satisfy a demand with N production units, at minimal cost
- Price decomposition
  - the coordinator sets a price λ_t
  - the units send their production $\mathbf{u}_t^{(i)}$
  - the coordinator compares total production and demand, and then updates the price
  - and so on...

Intuition of spatial decomposition

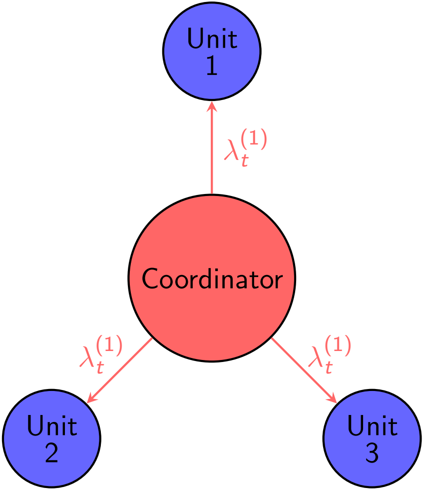

- Purpose: satisfy a demand with N production units, at minimal cost
- Price decomposition
  - the coordinator sets a price λ_t
  - the units send their production $\mathbf{u}_t^{(i)}$
  - the coordinator compares total production and demand, and then updates the price
  - and so on...

Intuition of spatial decomposition

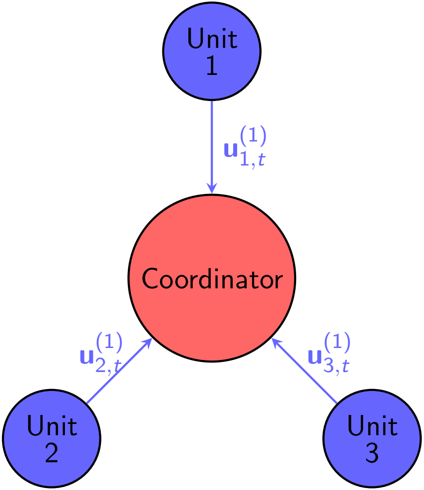

- Purpose: satisfy a demand with N production units, at minimal cost
- Price decomposition
  - the coordinator sets a price λ_t
  - the units send their production $\mathbf{u}^{(i)}_t$
  - the coordinator compares total production and demand, and then updates the price
  - and so on...

Intuition of spatial decomposition

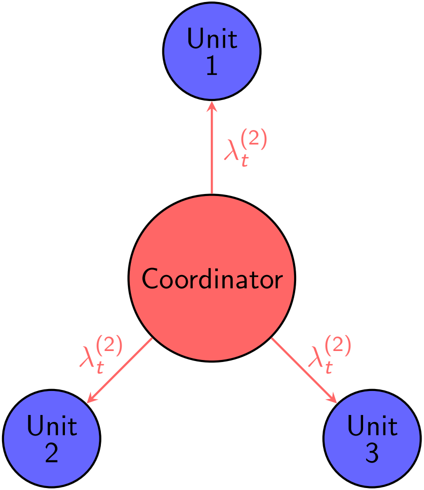

- Purpose: satisfy a demand with N production units, at minimal cost
- Price decomposition
  - the coordinator sets a price λ_t
  - the units send their production $\mathbf{u}_t^{(i)}$
  - the coordinator compares total production and demand, and then updates the price
  - and so on...

Intuition of spatial decomposition

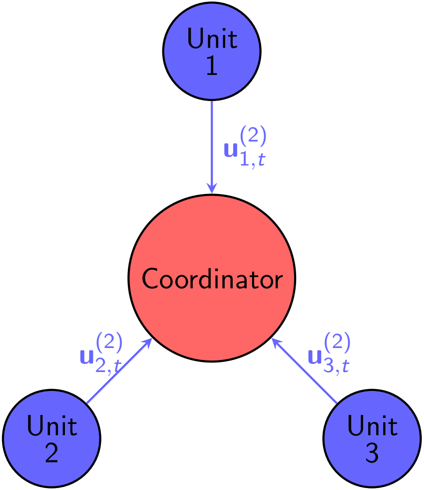

- Purpose: satisfy a demand with N production units, at minimal cost
- Price decomposition
  - the coordinator sets a price λ_t
  - the units send their production $\mathbf{u}^{(i)}_t$
  - the coordinator compares total production and demand, and then updates the price
  - and so on...

Price decomposition relies on dualization

$$\min_{u_i \in \mathcal{U}_i, i=1\ldots N} \sum_{i=1}^{N} J_i(u_i) \quad \text{subject to} \quad \sum_{i=1}^{N} \theta_i(u_i) - \theta = 0$$

1. Form the Lagrangian and assume that a saddle point exists

$$\max_{\lambda \in \mathcal{V}} \min_{u_i \in \mathcal{U}_i, i=1\ldots N} \sum_{i=1}^{N} \Big(J_i(u_i) + \left\langle \lambda \,, \theta_i(u_i) \right\rangle \Big) - \left\langle \lambda \,, \theta \right\rangle$$

2. Solve this problem by the dual gradient algorithm "à la Uzawa"

$$u_i^{(k+1)} \in \arg\min_{u_i \in \mathcal{U}_i} J_i(u_i) + \left\langle \lambda^{(k)} \,, \theta_i(u_i) \right\rangle \,, \quad i = 1 \ldots, N$$

$$\lambda^{(k+1)} = \lambda^{(k)} + \rho \Big(\sum_{i=1}^{N} \theta_i \Big(u_i^{(k+1)} \Big) - \theta \Big)$$

Price decomposition relies on dualization

$$\min_{u_i \in \mathcal{U}_i, i=1\dots N} \sum_{i=1}^{N} J_i(u_i) \quad \text{subject to} \quad \sum_{i=1}^{N} \theta_i(u_i) - \theta = 0$$

1. Form the Lagrangian and assume that a saddle point exists

$$\max_{\lambda \in \mathcal{V}} \min_{u_i \in \mathcal{U}_i, i=1\dots N} \sum_{i=1}^{N} \Big(J_i(u_i) + \left\langle \lambda , \theta_i(u_i) \right\rangle \Big) - \left\langle \lambda , \theta \right\rangle$$

2. Solve this problem by the dual gradient algorithm "à la Uzawa"

$$u_i^{(k+1)} \in \arg\min_{u_i \in \mathcal{U}_i} J_i(u_i) + \left\langle \lambda^{(k)} , \theta_i(u_i) \right\rangle , \quad i = 1 \dots , N$$

$$\lambda^{(k+1)} = \lambda^{(k)} + \rho \left(\sum_{i=1}^{N} \theta_i \left(u_i^{(k+1)} \right) - \theta \right)$$

Remarks on decomposition methods

- ▷ The theory is available for infinite dimensional Hilbert spaces, and thus applies in the stochastic framework, that is, when the $\mathcal{U}_i$ are spaces of random variables
- ▷ The minimization algorithm used for solving the subproblems is not specified in the decomposition process
- ▷ New variables $\lambda^{(k)}$ appear in the subproblems arising at iteration k of the optimization process

$$\min_{u_i \in \mathcal{U}_i} J_i(u_i) + \left\langle \lambda^{(k)} , \theta_i(u_i) \right\rangle$$

- ▷ These variables are fixed when solving the subproblems, and do not cause any difficulty, at least in the deterministic case

Price decomposition applies to various couplings

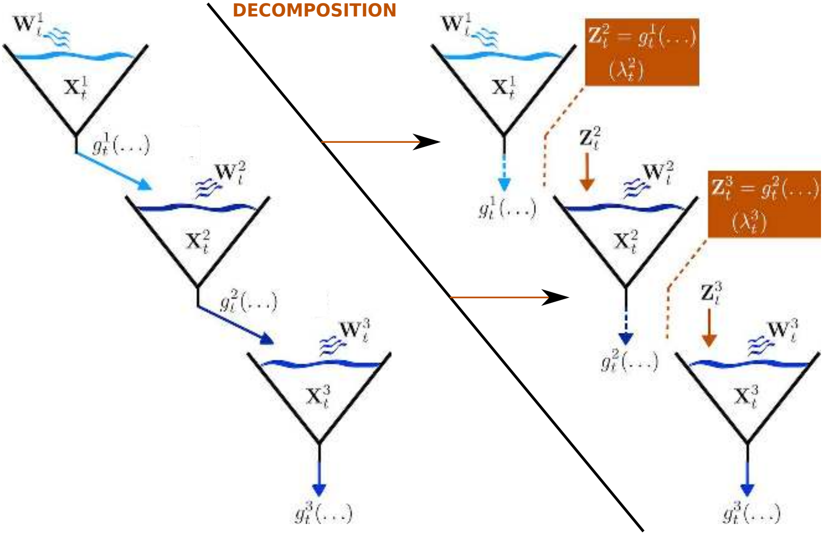

Stochastic optimal control (SOC) problem formulation

Consider the following SOC problem

$$\min_{\mathbf{u},\mathbf{x}} \ \mathbb{E}\Big(\sum_{i=1}^{N} \Big(\sum_{t=0}^{T-1} L_t^i(\mathbf{x}_t^i, \mathbf{u}_t^i, \mathbf{w}_{t+1}) + K^i(\mathbf{x}_T^i) \Big) \Big)$$

subject to the constraints

$$\begin{aligned}
\mathbf{x}_0^i &= f_{-1}^i(\mathbf{w}_0) \;, & i = 1 \dots N \\
\mathbf{x}_{t+1}^i &= f_t^i(\mathbf{x}_t^i, \mathbf{u}_t^i, \mathbf{w}_{t+1}) \;, & t = 0 \dots T-1 \;, \;\; i = 1 \dots N
\end{aligned}$$

$$\mathbf{u}_t^i \preceq \mathcal{F}_t = \sigma(\mathbf{w}_0, \dots, \mathbf{w}_t) \;, \;\; t = 0 \dots T-1 \;, \;\; i = 1 \dots N$$

$$\sum_{i=1}^{N} \theta_t^i(\mathbf{x}_t^i, \mathbf{u}_t^i) = 0 \;, \qquad t = 0 \dots T-1$$

Stochastic optimal control (SOC) problem formulation

Consider the following SOC problem

$$\min_{\mathbf{u},\mathbf{x}} \sum_{i=1}^{N} \left(\mathbb{E}\Big(\sum_{t=0}^{T-1} L_t^i(\mathbf{x}_t^i, \mathbf{u}_t^i, \mathbf{w}_{t+1}) + K^i(\mathbf{x}_T^i) \Big) \right)$$

subject to the constraints

$$\mathbf{x}_0^i = f_{-1}^i(\mathbf{w}_0) \; , \qquad i = 1 \ldots N$$

$$\mathbf{x}_{t+1}^i = f_t^i(\mathbf{x}_t^i, \mathbf{u}_t^i, \mathbf{w}_{t+1}) \; , \qquad t = 0 \ldots T-1 \; , \; i = 1 \ldots N$$

$$\mathbf{u}_t^i \preceq \mathcal{F}_t = \sigma(\mathbf{w}_0, \ldots, \mathbf{w}_t) \; , \; t = 0 \ldots T-1 \; , \; i = 1 \ldots N$$

$$\sum_{i=1}^{N} \theta_t^i(\mathbf{x}_t^i, \mathbf{u}_t^i) = 0 \; , \qquad t = 0 \ldots T-1$$

Dynamic Programming yields centralized controls

- As we want to solve this SOC problem using Dynamic Programming (**DP**), we suppose to be in the Markovian setting, that is, $\mathbf{w}_0, \dots, \mathbf{w}_T$ are a white noise
- The system is made of N interconnected subsystems, with the control $\mathbf{u}_t^i$ and the state $\mathbf{x}_t^i$ of subsystem i at time t
- The optimal control $\mathbf{u}_t^i$ of subsystem i is a function of the whole system state $\left(\mathbf{x}_t^1, \dots, \mathbf{x}_t^N\right)$
$$\mathbf{u}_t^i = \gamma_t^i\left(\mathbf{x}_t^1, \dots, \mathbf{x}_t^N\right)$$

Naive decomposition should lead to decentralized feedbacks

$$\mathbf{u}_t^i = \widehat{\gamma}_t^i(\mathbf{x}_t^i)$$

which are, in most cases, far from being optimal...

Straightforward decomposition of Dynamic Programming?

The crucial point is that the optimal feedback of a subsystem a priori depends on the state of all other subsystems, so that using a decomposition scheme by subsystems is not obvious...

As far as we have to deal with Dynamic Programming, the central concern for decomposition/coordination purpose boils down to

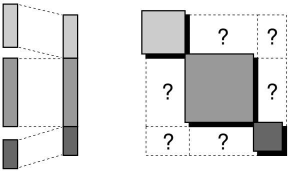

- ▷ how to decompose a feedback γ_t w.r.t. its domain $\mathbb{X}_t$ rather than its range $\mathbb{U}_t$?

  And the answer is
- ▷ impossible in the general case!

Price decomposition and Dynamic Programming

When applying price decomposition to the problem by dualizing the (almost sure) coupling constraint $\sum_i \theta_t^i(\mathbf{x}_t^i, \mathbf{u}_t^i) = 0$,
multipliers $\mathbf{\Lambda}_t^{(k)}$ appear in the subproblems arising at iteration k

$$\min_{\mathbf{u}^i, \mathbf{x}^i} \mathbb{E}\Big(\sum_t L_t^i(\mathbf{x}_t^i, \mathbf{u}_t^i, \mathbf{w}_{t+1}) + \mathbf{\Lambda}_t^{(k)} \cdot \theta_t^i(\mathbf{x}_t^i, \mathbf{u}_t^i) \Big)$$

- The variables $\mathbf{\Lambda}_t^{(k)}$ are fixed random variables, so that the random process $\mathbf{\Lambda}^{(k)}$ acts as an additional input noise in the subproblems
- But this process may be correlated in time, so that the white noise assumption has no reason to be fulfilled
- DP cannot be applied in a straightforward manner!

Question: how to handle the coordination instruments $\mathbf{\Lambda}_t^{(k)}$ to obtain (an approximation of) the overall optimum?

Price decomposition and Dynamic Programming

When applying price decomposition to the problem by dualizing the (almost sure) coupling constraint $\sum_i \theta_t^i(\mathbf{x}_t^i, \mathbf{u}_t^i) = 0$,
multipliers $\boldsymbol{\Lambda}_t^{(k)}$ appear in the subproblems arising at iteration k

$$\min_{\mathbf{u}^i, \mathbf{x}^i} \mathbb{E}\Big(\sum_t L_t^i(\mathbf{x}_t^i, \mathbf{u}_t^i, \mathbf{w}_{t+1}) + \boldsymbol{\Lambda}_t^{(k)} \cdot \theta_t^i(\mathbf{x}_t^i, \mathbf{u}_t^i) \Big)$$

- The variables $\boldsymbol{\Lambda}_t^{(k)}$ are fixed random variables, so that the random process $\boldsymbol{\Lambda}^{(k)}$ acts as an additional input noise in the subproblems
- But this process may be correlated in time, so that the white noise assumption has no reason to be fulfilled
- DP cannot be applied in a straightforward manner!

Question: how to handle the coordination instruments $\boldsymbol{\Lambda}_t^{(k)}$ to obtain (an approximation of) the overall optimum?

Optimization problem

The SOC problem under consideration reads

$$\min_{\mathbf{u},\mathbf{x}} \mathbb{E}\Big(\sum_{i=1}^{N} \Big(\sum_{t=0}^{T-1} L_t^i(\mathbf{x}_t^i, \mathbf{u}_t^i, \mathbf{w}_{t+1}) + K^i(\mathbf{x}_T^i) \Big) \Big) \tag{1a}$$

subject to dynamics constraints

$$\mathbf{x}_0^i = f_{-1}^i(\mathbf{w}_0) \tag{1b}$$

$$\mathbf{x}_{t+1}^i = f_t^i(\mathbf{x}_t^i, \mathbf{u}_t^i, \mathbf{w}_{t+1}) \tag{1c}$$

to measurability constraints:

$$\mathbf{u}_t^i \preceq \sigma(\mathbf{w}_0, \dots, \mathbf{w}_t) \tag{1d}$$

and to instantaneous coupling constraints

$$\sum_{i=1}^{N} \theta_t^i(\mathbf{x}_t^i, \mathbf{u}_t^i) = 0 \qquad \text{Constraints to be } \textbf{dualized} \tag{1e}$$

Assumptions

Assumption 1 (White noise)

Noises $\mathbf{w}_0, \ldots, \mathbf{w}_T$ are independent over time

Hence Dynamic Programming applies: there is no optimality loss to look after the controls $\mathbf{u}_t^i$ as functions of the state at time t

Assumption 2 (Constraint qualification)

A saddle point of the Lagrangian $\mathcal{L}$ exists

$$\mathcal{L}(\mathbf{x}, \mathbf{u}, \mathbf{\Lambda}) = \mathbb{E}\left(\sum_{i=1}^{N} \left(\sum_{t=0}^{T-1} L_t^i(\mathbf{x}_t^i, \mathbf{u}_t^i, \mathbf{w}_{t+1}) + K^i(\mathbf{x}_T^i) + \sum_{t=0}^{T-1} \mathbf{\Lambda}_t \cdot \theta_t^i(\mathbf{x}_t^i, \mathbf{u}_t^i) \right) \right)$$

where the $\mathbf{\Lambda}_t$ are $\sigma(\mathbf{w}_0, \ldots, \mathbf{w}_t)$-measurable random variables

Assumption 3 (Dual gradient algorithm)

Uzawa algorithm applies. . .

Uzawa algorithm

At iteration k of the algorithm,

1. Solve Subproblem i, $i = 1, \ldots, N$, with $\mathbf{\Lambda}^{(k)}$ fixed

$$\min_{\mathbf{u}^i, \mathbf{x}^i} \mathbb{E}\left(\sum_{t=0}^{T-1} \left(L_t^i(\mathbf{x}_t^i, \mathbf{u}_t^i, \mathbf{w}_{t+1}) + \mathbf{\Lambda}_t^{(k)} \cdot \theta_t^i(\mathbf{x}_t^i, \mathbf{u}_t^i) \right) + K^i(\mathbf{x}_T^i) \right)$$

subject to

$$\begin{aligned} \mathbf{x}_{t+1}^i &= f_t^i(\mathbf{x}_t^i, \mathbf{u}_t^i, \mathbf{w}_{t+1}) \\ \mathbf{u}_t^i &\preceq \sigma(\mathbf{w}_0, \ldots, \mathbf{w}_t) \end{aligned}$$

whose solution is denoted $\left(\mathbf{u}^{i,(k+1)}, \mathbf{x}^{i,(k+1)}\right)$

2. Update the multipliers $\mathbf{\Lambda}_t$

$$\mathbf{\Lambda}_t^{(k+1)} = \mathbf{\Lambda}_t^{(k)} + \rho_t \left(\sum_{i=1}^{N} \theta_t^i\left(\mathbf{x}_t^{i,(k+1)}, \mathbf{u}_t^{i,(k+1)}\right) \right)$$

Structure of a subproblem

- Subproblem i reads

$$\min_{\mathbf{u}^i,\mathbf{x}^i} \mathbb{E}\left(\sum_{t=0}^{T-1} \left(L_t^i(\mathbf{x}_t^i, \mathbf{u}_t^i, \mathbf{w}_{t+1}) + \boldsymbol{\Lambda}_t^{(k)} \cdot \theta_t^i(\mathbf{x}_t^i, \mathbf{u}_t^i) \right) \right)$$

subject to

$$\begin{aligned} \mathbf{x}_{t+1}^i &= f_t^i(\mathbf{x}_t^i, \mathbf{u}_t^i, \mathbf{w}_{t+1}) \\ \mathbf{u}_t^i &\preceq \sigma(\mathbf{w}_0, \ldots, \mathbf{w}_t) \end{aligned}$$

- Without some knowledge of the process $\boldsymbol{\Lambda}^{(k)}$ (we just know that $\boldsymbol{\Lambda}_t^{(k)} \preceq (\mathbf{w}_0, \ldots, \mathbf{w}_t)$), the informational state of this subproblem i at time t cannot be summarized by the physical state $\mathbf{x}_t^i$

We outline the main idea in DADP

- To overcome the difficulty induced by the term $\boldsymbol{\Lambda}_t^{(k)}$, we introduce a new adapted information process $\mathbf{y}^i = \big(\mathbf{y}_0^i, \ldots, \mathbf{y}_{T-1}^i\big)$ for Subsystem i
- at each time t, the random variable $\mathbf{y}_t^i$ is measurable w.r.t. the past noises $\big(\mathbf{w}_0, \ldots, \mathbf{w}_t\big)$
- The core idea is to replace the multiplier $\boldsymbol{\Lambda}_t^{(k)}$ at iteration k by its conditional expectation $\mathbb{E}(\boldsymbol{\Lambda}_t^{(k)} \mid \mathbf{y}_t^i)$
- (More on the interpretation later)

Note that we require that the information process is not influenced by controls

We can now approximate Subproblem i

- Using this idea, we replace Subproblem i by

$$\min_{\mathbf{u}^i,\mathbf{x}^i} \mathbb{E}\Big(\sum_{t=0}^{T-1} \Big(L_t^i(\mathbf{x}_t^i, \mathbf{u}_t^i, \mathbf{w}_{t+1}) + \mathbb{E}(\mathbf{\Lambda}_t^{(k)} \mid \mathbf{y}_t^i) \cdot \theta_t^i(\mathbf{x}_t^i, \mathbf{u}_t^i) \Big) + K^i(\mathbf{x}_T^i) \Big)$$

subject to

$$\mathbf{x}_{t+1}^i = f_t^i(\mathbf{x}_t^i, \mathbf{u}_t^i, \mathbf{w}_{t+1})$$
$$\mathbf{u}_t^i \preceq \sigma(\mathbf{w}_0, \ldots, \mathbf{w}_t)$$

- The conditional expectation $\mathbb{E}(\mathbf{\Lambda}_t^{(k)} \mid \mathbf{y}_t^i)$ is an (updated) function of the variable $\mathbf{y}_t^i$,
- so that Subproblem i involves the two noises processes $\mathbf{w}$ and $\mathbf{y}^i$

If $\mathbf{y}^i$ follows a dynamical equation, DP applies

We obtain a Dynamic Programming equation by subsystem

Assuming a non-controlled dynamics $\mathbf{y}^i_{t+1} = h^i_t\big(\mathbf{y}^i_t, \mathbf{w}_{t+1}\big)$ for the information process $\mathbf{y}^i$, the DP equation writes

$$
\begin{aligned}
V^i_T(x,y) &= K^i(x) \\
V^i_t(x,y) &= \min_u \mathbb{E}\Big(L^i_t(x,u,\mathbf{w}_{t+1}) \\
&\qquad\qquad + \mathbb{E}\big(\mathbf{\Lambda}^{(k)}_t \mid \mathbf{y}^i_t = y\big) \cdot \theta^i_t(x,u) \\
&\qquad\qquad\qquad + V^i_{t+1}\big(\mathbf{x}^i_{t+1}, \mathbf{y}^i_{t+1}\big)\Big)
\end{aligned}
$$

subject to the dynamics

$$
\begin{aligned}
\mathbf{x}^i_{t+1} &= f^i_t(x,u,\mathbf{w}_{t+1}) \\
\mathbf{y}^i_{t+1} &= h^i_t(y,\mathbf{w}_{t+1})
\end{aligned}
$$

DADP displays three interpretations

- DADP as an approximation of the optimal multiplier

$$\lambda_t \quad \rightsquigarrow \quad \mathbb{E}\big(\lambda_t \mid \mathbf{y}_t\big)$$

- DADP as a decision-rule approach in the dual

$$\max_{\lambda} \min_{\mathbf{u}} L\big(\lambda, \mathbf{u}\big) \quad \rightsquigarrow \quad \max_{\lambda_t \preceq \mathbf{y}_t} \min_{\mathbf{u}} L\big(\lambda, \mathbf{u}\big)$$

- DADP as a constraint relaxation

$$\sum_{i=1}^{n} \theta_t^i\big(\mathbf{u}_t^i\big) = 0 \quad \rightsquigarrow \quad \mathbb{E}\bigg(\sum_{i=1}^{n} \theta_t^i\big(\mathbf{u}_t^i\big) \;\bigg|\; \mathbf{y}_t\bigg) = 0$$

A bunch of practical questions remains open

- ⋆ How to choose the information variables $\mathbf{y}_t^i$?
  - ▷ Perfect memory: $\mathbf{y}_t^i = \big(\mathbf{w}_0, \ldots, \mathbf{w}_t\big)$
  - ▷ Minimal information: $\mathbf{y}_t^i \equiv \text{cste}$
  - ▷ Static information: $\mathbf{y}_t^i = h_t^i\big(\mathbf{w}_t\big)$
  - ▷ Dynamic information: $\mathbf{y}_{t+1}^i = h_t^i\big(\mathbf{y}_t^i, \mathbf{w}_{t+1}\big)$
- ⋆ How to obtain a feasible solution from the relaxed problem?
  - ▷ Use an appropriate heuristic!
- ⋆ How to accelerate the gradient algorithm?
  - ▷ Augmented Lagrangian
  - ▷ More sophisticated gradient methods
- ⋆ How to handle more complex structures than the flower model?

We consider 3 dams in a row, amenable to DP

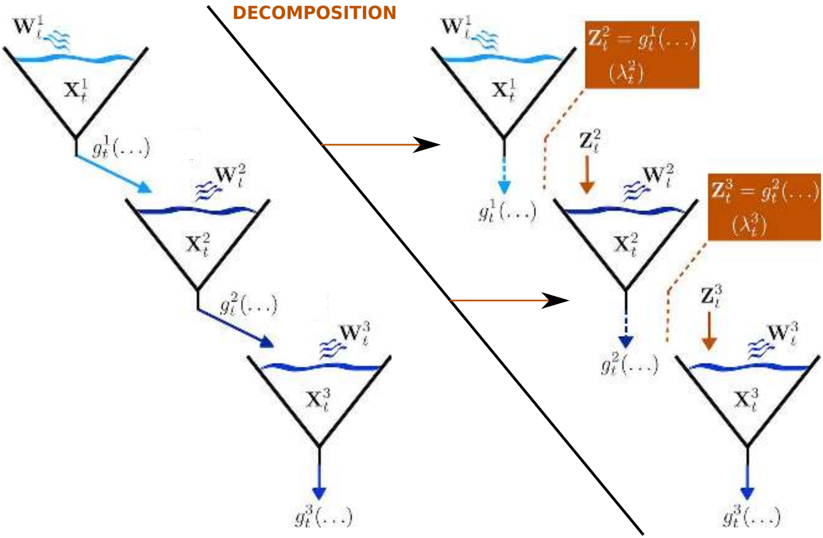

Problem specification

- ▷ We consider a 3 dam problem, over 12 time steps
- ▷ We relax each constraint with a given information process $\mathbf{y}^i$
- ▷ All random variable are discrete (noise, control, state)

- ▷ We test the following information processes

Constant information: equivalent to replace the a.s. constraint by an expected constraint

Part of noise: the information process is the inflow of the above dam $\mathbf{Y}_t^i = \mathbf{w}_t^{i-1}$

Phantom state: the information process mimicks the optimal trajectory of the state of the first dam (by statistical regression over the known optimal trajectory in this case)

Numerical results are encouraging

	DADP - $\mathbb{E}$	DADP - $\mathbf{w}^{i-1}$	DADP - dyn.	DP
Nb of it.	165	170	25	1
Time (min)	2	3	67	41
Lower Bound	-1.386×10^6	-1.379×10^6	-1.373×10^6	
Final Value	-1.335×10^6	-1.321×10^6	-1.344×10^6	-1.366×10^6
Loss	-2.3%	-3.3%	-1.6%	ref.

$\leadsto$ *PhD thesis of J.-C. Alais*

Summing up DADP

- Choose an information process $\mathbf{y}$ following $\mathbf{y}_{t+1} = \widetilde{f}_t(\mathbf{y}_t, \mathbf{w}_{t+1})$
- Relax the almost sure coupling constraint into a conditional expectation
- Then apply a price decomposition scheme to the relaxed problem
- The subproblems can be solved by dynamic programming with the modest state $(\mathbf{x}^i_t, \mathbf{y}_t)$
- In the theoretical part, we give
  - a consistency result (family of information process)
  - a convergence result (fixed information process)
  - conditions for the existence of multiplier

What are the issues to consider?

- ▷ We treat the coupling constraints in a stochastic optimization problem by duality methods
- ▷ Uzawa algorithm is a dual method which is naturally described in an Hilbert space, but we cannot guarantee the existence of an optimal multiplier in the space $\mathrm{L}^2\big(\Omega, \mathcal{F}, \mathbb{P}; \mathbb{R}^n\big)$!
- ▷ Consequently, we extend the algorithm to the non-reflexive Banach space $\mathrm{L}^\infty\big(\Omega, \mathcal{F}, \mathbb{P}; \mathbb{R}^n\big)$, by giving a set of conditions ensuring the existence of a $\mathrm{L}^1\big(\Omega, \mathcal{F}, \mathbb{P}; \mathbb{R}^n\big)$ optimal multiplier,
  and by providing a convergence result of the algorithm

- ▷ We also have to deal with the approximation induced by the information variable: we give an epi-convergence result related to such an approximation

$\rightsquigarrow$ *PhD thesis of V. Leclère*

Abstract formulation of the problem

We consider the following abstract optimization problem

$$(\mathcal{P}) \qquad \min_{\mathbf{u}\in\mathcal{U}^{\mathrm{ad}}} J(\mathbf{u}) \quad \text{s.t.} \quad \Theta(\mathbf{u}) \in -C$$

where $\mathcal{U}$ and $\mathcal{V}$ are two Banach spaces, and

- $J : \mathcal{U} \to \overline{\mathbb{R}}$ is the objective function
- $\mathcal{U}^{\mathrm{ad}}$ is the admissible set
- $\Theta : \mathcal{U} \to \mathcal{V}$ is the constraint function to be dualized
- $C \subset \mathcal{V}$ is the cone of constraint

Let $\mathcal{U}^{\Theta} = \{\mathbf{u} \in \mathcal{U},\ \Theta(\mathbf{u}) \in -C\}$ be the associated constraint set

Here, $\mathcal{U}$ is a space of random variables, and J is defined by

$$J(\mathbf{u}) = \mathbb{E}\big(j(\mathbf{u}, \mathbf{w})\big)$$

The relationship with Problem (1) is almost straightforward. . .

Standard duality in L^2 spaces (I)

Assume that $\mathcal{U} = \mathrm{L}^2(\Omega, \mathcal{F}, \mathbb{P}; \mathbb{R}^n)$ and $\mathcal{V} = \mathrm{L}^2(\Omega, \mathcal{F}, \mathbb{P}; \mathbb{R}^m)$

The standard sufficient constraint qualification condition

$$0 \in \mathrm{ri}\Big(\Theta\big(\mathcal{U}^{\mathrm{ad}} \cap \mathrm{dom}(J)\big) + C\Big)$$

is scarcely satisfied in such a stochastic setting

Proposition 1

If the σ-algebra $\mathcal{F}$ is not finite modulo $\mathbb{P}$,[a]
then for any subset $U^{\mathrm{ad}} \subset \mathbb{R}^n$ that is not an affine subspace, the set

$$\mathcal{U}^{\mathrm{ad}} = \Big\{\mathbf{u} \in \mathrm{L}^p(\Omega, \mathcal{F}, \mathbb{P}; \mathbb{R}^n) \mid \mathbf{u} \in U^{\mathrm{ad}} \quad \mathbb{P}-a.s.\Big\}$$

has an empty relative interior in L^p, for any $p < +\infty$

[a] If the σ-algebra is finite modulo $\mathbb{P}$, $\mathcal{U}$ and $\mathcal{V}$ are finite dimensional spaces

Standard duality in L^2 spaces (II)

Consider the following optimization problem:

$$\begin{aligned} \inf_{u_0,\mathbf{u}_1} \quad & u_0^2 + \mathbb{E}\big((\mathbf{u}_1 + \alpha)^2\big) \\ \text{s.t.} \quad & u_0 \geq a \\ & \mathbf{u}_1 \geq 0 \\ & u_0 - \mathbf{u}_1 \geq \mathbf{w} \qquad \text{to be dualized} \end{aligned}$$

where $\mathbf{w}$ is a random variable uniform on $[1, 2]$

For $a < 2$, we can construct a maximizing sequence of multipliers for the dual problem that does not converge in L^2.
(We are in the so-called *non relatively complete recourse* case, that is, the case where the constraints on $\mathbf{u}_1$ induce a stronger constraint on u_0)

An optimal multiplier is available in $(\mathrm{L}^\infty)^\star$...

Standard duality in L^2 spaces (II)

Consider the following optimization problem:

$$\begin{aligned} \inf_{u_0,\mathbf{u}_1} \quad & u_0^2 + \mathbb{E}\left((\mathbf{u}_1 + \alpha)^2\right) \\ \text{s.t.} \quad & u_0 \geq a \\ & \mathbf{u}_1 \geq 0 \\ & u_0 - \mathbf{u}_1 \geq \mathbf{w} \qquad \text{to be dualized} \end{aligned}$$

where $\mathbf{w}$ is a random variable uniform on $[1, 2]$

For $a < 2$, we can construct a maximizing sequence of multipliers for the dual problem that does not converge in L^2.
(We are in the so-called *non relatively complete recourse* case, that is, the case where the constraints on $\mathbf{u}_1$ induce a stronger constraint on u_0)

An optimal multiplier is available in $\left(\mathrm{L}^\infty\right)^\star$...

Constraint qualification in $(\mathrm{L}^\infty, \mathrm{L}^1)$

From now on, we assume that

$$\mathcal{U} = \mathrm{L}^\infty(\Omega, \mathcal{F}, \mathbb{P}; \mathbb{R}^n)$$
$$\mathcal{V} = \mathrm{L}^\infty(\Omega, \mathcal{F}, \mathbb{P}; \mathbb{R}^m)$$
$$C = \{0\}$$

where the σ-algebra $\mathcal{F}$ is not finite modulo $\mathbb{P}$

We consider the pairing $(\mathrm{L}^\infty, \mathrm{L}^1)$ with the following topologies:

- ▷ $\sigma\big(\mathrm{L}^\infty, \mathrm{L}^1\big)$: weak⋆ topology on L^∞ (coarsest topology such that all the L^1-linear forms are continuous),
- ▷ $\tau\big(\mathrm{L}^\infty, \mathrm{L}^1\big)$: Mackey-topology on L^∞ (finest topology such that the continuous linear forms are only the L^1-linear forms)

Weak* closedness of linear subspaces of L^∞

Proposition 2

Let $\Theta : \mathrm{L}^\infty(\Omega, \mathcal{F}, \mathbb{P}; \mathbb{R}^n) \to \mathrm{L}^\infty(\Omega, \mathcal{F}, \mathbb{P}; \mathbb{R}^m)$ *be a linear operator, and assume that there exists a linear operator* $\Theta^\dagger : \mathrm{L}^1(\Omega, \mathcal{F}, \mathbb{P}; \mathbb{R}^m) \to \mathrm{L}^1(\Omega, \mathcal{F}, \mathbb{P}; \mathbb{R}^n)$ *such that:*

$$\langle \mathbf{v} , \Theta(\mathbf{u}) \rangle = \langle \Theta^\dagger(\mathbf{v}) , \mathbf{u} \rangle , \quad \forall \mathbf{u}, \ \forall \mathbf{v}$$

Then the linear operator Θ *is weak** *continuous*

Applications

- $\Theta(\mathbf{u}) = \mathbf{u} - \mathbb{E}(\mathbf{u} \mid \mathcal{B})$: non-anticipativity constraints,
- $\Theta(\mathbf{u}) = A\mathbf{u}$ with $A \in \mathcal{M}_{m,n}(\mathbb{R})$: finite number of constraints

A duality theorem

$$(\mathcal{P}) \qquad \min_{\mathbf{u}\in\mathcal{U}} J(\mathbf{u}) \quad \text{s.t.} \quad \Theta(\mathbf{u}) = 0$$

with $J(\mathbf{u}) = \mathbb{E}\big(j(\mathbf{u},\mathbf{w})\big)$

Theorem 1

Assume that j is a convex normal integrand, that J is continuous in the Mackey topology at some point $\mathbf{u}_0$ such that $\Theta(\mathbf{u}_0) = 0$, and that Θ is weak* continuous on $\mathrm{L}^\infty\big(\Omega,\mathcal{F},\mathbb{P};\mathbb{R}^n\big)$
Then, $\mathbf{u}^\star \in \mathcal{U}$ is an optimal solution of Problem $\big(\mathcal{P}\big)$ if and only if there exists $\lambda^\star \in \mathrm{L}^1\big(\Omega,\mathcal{F},\mathbb{P};\mathbb{R}^m\big)$ such that

- ▷ $\mathbf{u}^\star \in \arg\min_{\mathbf{u}\in\mathcal{U}} \mathbb{E}\Big(j(\mathbf{u},\mathbf{w}) + \lambda^\star \cdot \Theta(\mathbf{u})\Big)$
- ▷ $\Theta(\mathbf{u}^\star) = 0$

Extension of a result given by R. Wets for non-anticipativity constraints

Uzawa algorithm

$$(\mathcal{P}) \qquad \min_{\mathbf{u}\in\mathcal{U}} J(\mathbf{u}) \quad \text{s.t.} \quad \Theta(\mathbf{u}) = 0$$

with $J(\mathbf{u}) = \mathbb{E}\big(j(\mathbf{u},\mathbf{w})\big)$

The standard Uzawa algorithm

$$\mathbf{u}^{(k+1)} \in \operatorname*{arg\,min}_{\mathbf{u}\in\mathcal{U}^{\mathrm{ad}}} J(\mathbf{u}) + \left\langle \lambda^{(k)} , \Theta(\mathbf{u}) \right\rangle$$

$$\lambda^{(k+1)} = \lambda^{(k)} + \rho\, \Theta\big(\mathbf{u}^{(k+1)}\big)$$

makes sense with in the L^∞ setting, that is, the minimization problem is well-posed and the update formula is valid one

Note that all the multipliers $\lambda^{(k)}$ *belong to* $\mathrm{L}^\infty\big(\Omega,\mathcal{F},\mathbb{P};\mathbb{R}^m\big)$*,*
as soon as the initial multiplier $\lambda^{(0)} \in \mathrm{L}^\infty\big(\Omega,\mathcal{F},\mathbb{P};\mathbb{R}^m\big)$

Convergence result

Theorem 2

Assume that

1. $J : \mathcal{U} \to \overline{\mathbb{R}}$ *is proper, weak* l.s.c., differentiable and* a*-convex*
2. $\Theta : \mathcal{U} \to \mathcal{V}$ *is affine, weak* continuous and* κ*-Lipschitz*
3. $\mathcal{U}^{\mathrm{ad}}$ *is weak* closed and convex,*
4. *an admissible* $\mathbf{u}_0 \in \mathrm{dom}\, J \cap \Theta^{-1}(0) \cap \mathcal{U}^{\mathrm{ad}}$ *exists*
5. *an optimal* L^1*-multiplier to the constraint* $\Theta\big(\mathbf{u}\big) = 0$ *exists*
6. *the step* ρ *is such that* $0 < \rho < \frac{2a}{\kappa}$

Then, there exists a subsequence $\big\{\mathbf{u}^{(n_k)}\big\}_{k \in \mathbb{N}}$ *of the sequence generated by the Uzawa algorithm converging in* L^∞ *toward the optimal solution* $\mathbf{u}^\star$ *of the primal problem*

Remarks about these results

- ▷ The result is not as good as expected (global convergence)
- ▷ Improvements and extensions (inequality constraint) needed
- ▷ The Mackey-continuity assumption forbids the use of bounds
  - ▷ In order to deal with almost sure bound constraints, we can turn towards the work of R.T. Rockafellar and R. J-B Wets
  - ▷ In a series of 4 papers (stochastic convex programming), they have detailed the duality theory on two-stage and multistage problems, with the focus on non-anticipativity constraints
  - ▷ These papers require
    - a strict feasability assumption
    - a relatively complete recourse assumption

Relaxed problems

Following the interpretation of DADP in terms of a relaxation of the original problem, and given a sequence $\{\mathcal{F}_n\}_{n\in\mathbb{N}}$ of subfields of the σ-field $\mathcal{F}$, we replace the abstract problem

$$(\mathcal{P}) \qquad \min_{\mathbf{u}\in\mathcal{U}} J(\mathbf{u}) \quad \text{s.t.} \quad \Theta(\mathbf{u}) = 0$$

by the sequence of approximated problems:

$$(\mathcal{P}_n) \qquad \min_{\mathbf{u}\in\mathcal{U}} J(\mathbf{u}) \quad \text{s.t.} \quad \mathbb{E}\big(\Theta(\mathbf{u}) \bigm| \mathcal{F}_n\big) = 0$$

We assume the Kudo convergence of $\{\mathcal{F}_n\}_{n\in\mathbb{N}}$ toward $\mathcal{F}$:

$$\mathcal{F}_n \longrightarrow \mathcal{F} \iff \forall \mathbf{x} \in \mathrm{L}^1(\Omega, \mathcal{F}, \mathbb{P}; \mathbb{R}),\ \mathbb{E}\big(\mathbf{x} \bigm| \mathcal{F}_n\big) \xrightarrow{\mathrm{L}^1} \mathbb{E}\big(\mathbf{x} \bigm| \mathcal{F}\big)$$

Convergence result

Theorem 3

Assume that

- ▷ $\mathcal{U}$ *is a topological space*
- ▷ $\mathcal{V} = \mathrm{L}^p(\Omega, \mathcal{F}, \mathbb{P}; \mathbb{R}^m)$ *with* $p \in [1, +\infty)$
- ▷ J *and* Θ *are continuous operators*
- ▷ $\{\mathcal{F}_n\}_{n\in\mathbb{N}}$ *Kudo converges toward* $\mathcal{F}$

Then the sequence $\{\widetilde{J}_n\}_{n\in\mathbb{N}}$ *epi-converges toward* $\widetilde{J}$*, with*

$$\widetilde{J}_n = \begin{cases} J(\mathbf{u}) & \text{if } \mathbf{u} \text{ satisfies the constraints of } (\mathcal{P}_n) \\ +\infty & \text{otherwise} \end{cases}$$

Conclusion

- DADP method allows to tackle
  large-scale stochastic optimal control problems,
  such as those found in energy management
- A host of theoretical and practical questions remains open
- We would like to test DADP on (smart) grids, extending the works
  on “flower models” (Unit Commitment problem) and
  on “chained models” (hydraulic valley management)
  to “network models” (grids)