Sequential Decision Models under Uncertainty

Extended from Chapter 6 of
Sustainable Management of Natural Resources. Mathematical Models and Methods
by Luc DOYEN and Michel DE LARA

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

École des Ponts ParisTech

August 25, 2014

Outline of the presentation

1. Dynamical control systems under uncertainty
2. Scenarios support a priori/off-line information
3. On-line information feeds policies
4. Summary

Outline of the presentation

Outline of the presentation

Uncertainty is pervasive in natural resources management

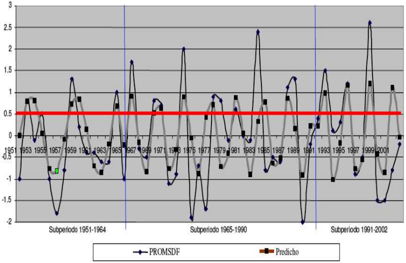

- Environmental uncertainties (*El Niño*)
- Habitats changes, mortality, natality
- Scientific uncertainties (structure of trophic networks, ecosystem services)

We plug incertain variables into the carbon cycle model "à la Nordhaus"

- Economic production $Q(t)$

$$Q(t+1) = \Big(1 + \overbrace{g\big(w_e(t)\big)}^{\text{economic growth}}\Big) Q(t)$$

- CO_2 concentration $M(t)$

$$M(t+1) = M(t) - \delta(M(t) - M_{-\infty}) + \underbrace{\alpha(w_p(t))}_{\text{physics}} \overbrace{\mathtt{Emiss}\big(Q(t), w_z(t)\big)}^{\text{technologies}} \big(1 - a(t)\big)$$

- Vector of uncertainties $w(t) = (w_e(t), w_p(t), w_z(t))$ on
  - economic growth
  - technologies
  - climate dynamics

Uncertainties transpire in epidemiological models

- Basic variables and parameters are
  - time t, measured in weeks
  - M_t, the abundance of infected mosquitos (Aedes Aegypti adultos)
  - H_t, the abundance of infected humans
  - $\Delta\mu_t^M$, the additional mortality rate of mosquitos, a control variable
  - $\overline{M}$, $\overline{H}$, f^H, f^M, μ^M and μ^H, parameters
- The controlled dynamics is

$$
\begin{aligned}
M_{t+1} &= f^H H_t(\overline{M} - M_t) - (\mu^M + \Delta\mu_t^M)M_t \\
H_{t+1} &= f^M M_t(\overline{H} - H_t) - \mu^H H_t
\end{aligned}
$$

- Scientific literature provides bounds for
  - disease transmission rates f^H and f^M
  - mortality rate of mosquitos μ^M

Uncertainties abound in population models

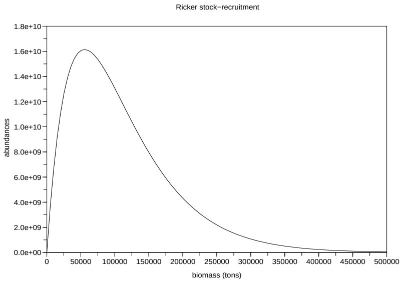

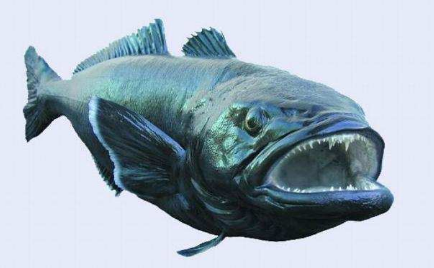

- Stock-recruitment relationship condenses, in one function, complex mechanisms of birth, dispersion, predation, habitats, physical conditions, etc.
- Natural mortality (deseases, predation) between age-classes is poorly known

We plug incertain variables into the harvested age-class model

$$
\begin{aligned}
N_1(t+1) &= \mathtt{S/R}\Big(\mathtt{SSB}(N(t)), \underbrace{w(t)}_{\text{birth mortality, etc.}}\Big) \quad \text{recruitment} \\
N_2(t+1) &= e^{-(M_1+\lambda(t)F_1)} N_1(t) \\
\vdots &= \vdots \\
N_a(t+1) &= e^{-(\overbrace{M_{a-1}}^{\text{mortality}}+\lambda(t)F_{a-1})} N_{a-1}(t), \quad a = 2, \ldots, A-1 \\
N_A(t+1) &= e^{-(M_{A-1}+\lambda(t)F_{A-1})} N_{A-1}(t) + \pi e^{-(M_A+\lambda(t)F_A)} N_A(t)
\end{aligned}
$$

Outline of the presentation

Uncertainty variables are new input variables

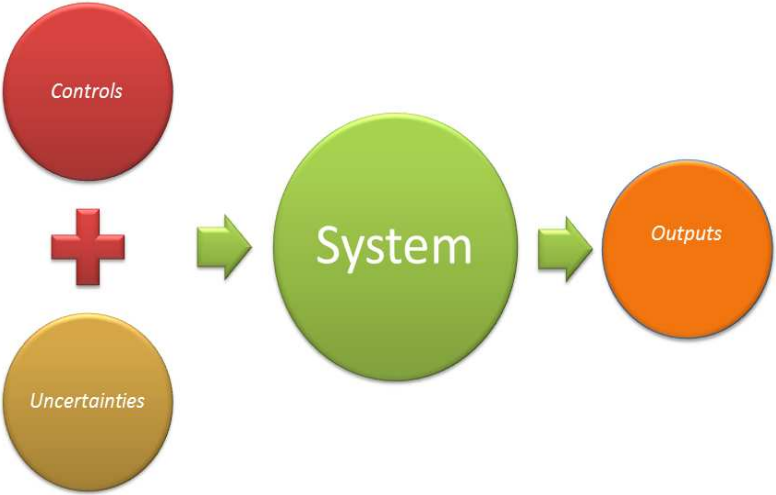

Input control variables are in the hands of the decision-maker at successive time periods

Control variables $u(t) \in \mathbb{U}$

The decision-maker can choose the values of control variables $u(t)$ at any time within given bounds

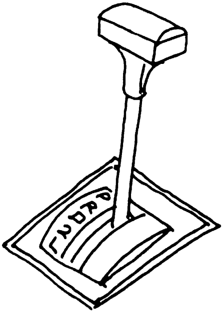

- at successive time periods
  - annual catches
  - years, months:
    starting of energy units like nuclear plants
  - weeks, days, intra-day: starting of hydropower units
- within given bounds
  - fishing quotas
  - turbined capacity

Input uncertain variables are out of the control of the decision-maker

Uncertain variables $w(t) \in \mathbb{W}$ are variables

- that take more than one single value (else they are deterministic)
- and over which the decision-maker (DM) has no control whatsoever

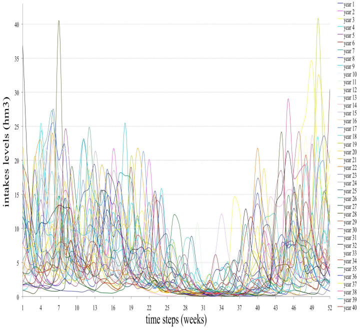

- Stationary parameters:
  unitary cost of CO_2 emissions
- Trends or seasonal effects:
  energy consumption pathway, mean temperatures, mean prices
- Stochastic processes:
  rain inputs in a dam, energy demand, prices
- Else (set membership):
  costs of climate change damage,
  water inflows in a dam

Uncertainty variables are new input variables in a discrete-time nonlinear state-control system

A specific output is distinguished, and is labeled "state" (more on this later), when the system may be written

$$x(t+1) = \mathtt{Dyn}\big(t, x(t), u(t), w(t)\big)\ , \quad t \in \mathbb{T} = \{t_0, t_0+1, \ldots, T-1\}$$

- time $t \in \overline{\mathbb{T}} = \{t_0, t_0+1, \ldots, T-1, T\} \subset \mathbb{N}$ (the time period $[t, t+1[$ may be a year, a month, etc.)
- state $x(t) \in \mathbb{X} := \mathbb{R}^n$ (biomasses, abundances, etc.)
- control $u(t) \in \mathbb{U} := \mathbb{R}^p$ (catches or harvesting effort)
- uncertainty $w(t) \in \mathbb{W} := \mathbb{R}^q$ (recruitment or mortality uncertainties, climate fluctuations or trends, etc.)
- dynamics $\mathtt{Dyn}$ maps $\mathbb{T} \times \mathbb{X} \times \mathbb{U} \times \mathbb{W}$ into $\mathbb{X}$ (biomass model, age-class model, economic model)

What have we covered so far?

Uncertainty variables are new input variables

$$x(t+1) = \texttt{Dyn}\Big(t, x(t), u(t), \underbrace{w(t)}_{\text{uncertainty}}\Big)$$

- The future state $x(t+1)$ is no longer predictable
- because of the uncertain term $w(t)$,
- but the current state $x(t)$ carries information relevant for decision-making,
- and we shed light on the notion of policy

Summary

- Control variables are defined rather unambiguously:
  the DM can select their values at any time within given sets
- The distinction between input and output variables is relative to a system:
  for two interconnected dams, the water release from the upper to the lower dam can be "seen" as an input to the lower dam or as a control variable for the two-dams system
- In various examples of natural resources management,
  we have seen so-called uncertain variables
- Uncertain variables are variables
  - which take more than one single value (else they are deterministic)
  - and over which the decision-makers have no control whatsoever
- Uncertain and control variables combine in a dynamical model

Outline of the presentation

Outline of the presentation

1. Dynamical control systems under uncertainty
   - Examples of uncertainties in dynamical systems
   - Uncertainty variables are new input variables
2. Scenarios support a priori/off-line information
   - Scenarios are temporal sequence of uncertainties
   - A priori / off-line information
3. On-line information feeds policies
   - The concept of policy
   - State and control solution maps
4. Summary

Water inflows historical scenarios

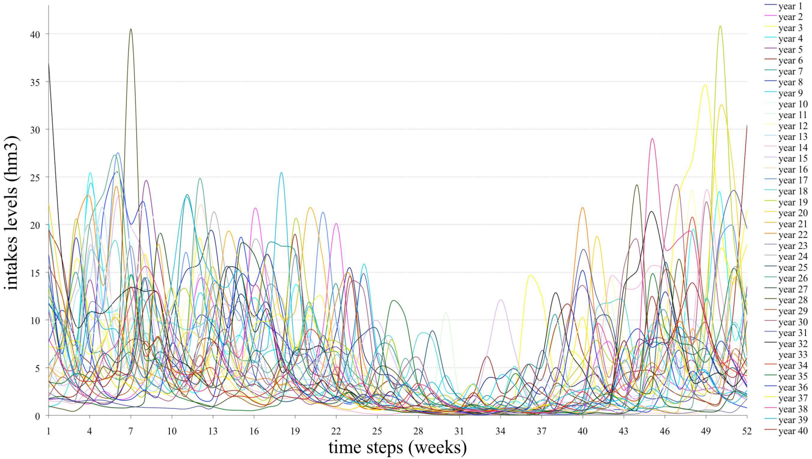

We call scenario a temporal sequence of uncertainties

Scenarios are special cases of "states of Nature"

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

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

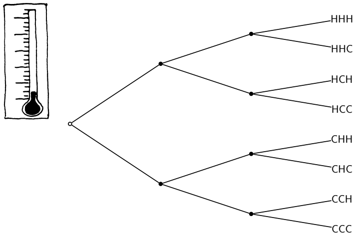

El tiempo se bifurca perpetuamente hacia innumerables futuros
(Jorge Luis Borges, *El jardín de senderos que se bifurcan*)

Beware! Scenario holds a different meaning in other scientific communities

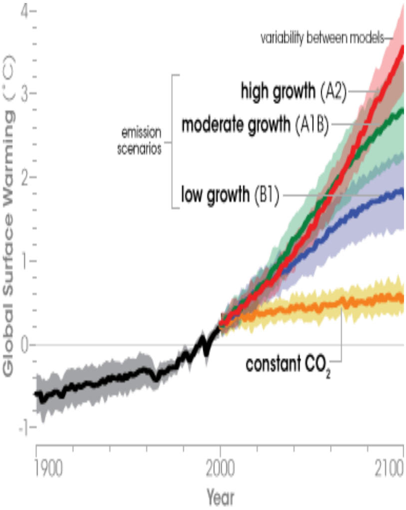

- In practice, what modelers call a "scenario" is a mixture of
  - a sequence of uncertain variables (also called a pathway, a chronicle)
  - a policy `Pol`
  - and even a static or dynamical model
- In what follows

  scenario = pathway = chronicle

Outline of the presentation

1. Dynamical control systems under uncertainty
   - Examples of uncertainties in dynamical systems
   - Uncertainty variables are new input variables
2. Scenarios support a priori/off-line information
   - Scenarios are temporal sequence of uncertainties
   - A priori / off-line information
3. On-line information feeds policies
   - The concept of policy
   - State and control solution maps
4. Summary

Probabilistic and set-membership approaches are ways to translate a priori / off-line information as illustrated in nuclear accidents prevention

- Three Mile Island accident:
  before the fact, the core meltdown was considered as excluded
- Nuclear accidents with probability per reactor per year
  - between 10^{-6} and 10^{-4} are considered as hypothetical,
  - whereas below 10^{-6} they are not envisaged
- Fukushima nuclear plants had a 10^{-9} nuclear accident probability per reactor per year

Choosing a set of scenarios is excluding "things we do not know we don't know"

Reports that say that something hasn't happened are always interesting to me, because as we know, there are known knowns; there are things we know we know. We also know there are known unknowns; that is to say we know there are some things we do not know. But there are also unknown unknowns – the ones we don't know we don't know. And if one looks throughout the history of our country and other free countries, it is the latter category that tend to be the difficult ones.

Donald Rumsfeld, former United States Secretary of Defense. From Department of Defense news briefing, February 12, 2002

In the stochastic approach, the set of scenarios is equipped with a known probability

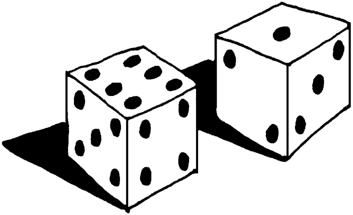

A priori information on the scenarios may be probabilistic

- A probability distribution $\mathbb{P}$ on Ω

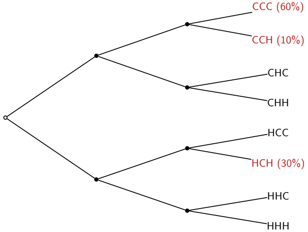

- In practice, one often assumes that the components $\big(w(t_0), \ldots, w(T-1)\big)$ form
  - an independent and identically distributed sequence
  - a Markov chain, a time series, etc.

Water inflows in a dam

Water inflows in a dam may be modelled as time series (ARMA, etc.)

Probabilistic assumptions and expected value

- The domain of scenarios $\Omega = \mathbb{W}^{T+1-t_0} = \mathbb{R}^q \times \cdots \times \mathbb{R}^q$ is equipped with the σ-field $\mathcal{F} = \bigotimes_{t=t_0}^{T} \mathcal{B}(\mathbb{R}^q)$ and a probability $\mathbb{P}$
- The sequences $w(\cdot) = (w(t_0), w(t_0+1), \ldots, w(T-1), w(T))$ now become the primitive random variables
- The notation $\mathbb{E}_{\mathbb{P}}$ refers to the mathematical expectation over Ω under probability $\mathbb{P}$

$$\mathbb{E}[A(w(\cdot))] = \sum_{w(\cdot)\in\Omega} \mathbb{P}\{w(\cdot)\} A(w(\cdot))$$

- The expectation operator $\mathbb{E}_{\mathbb{P}}$ enjoys linearity in the $(+, \times)$ algebra:

$$\mathbb{E}_{\mathbb{P}}(A+B) = \mathbb{E}_{\mathbb{P}}(A) + \mathbb{E}_{\mathbb{P}}(B)$$

- The random variables $(w(t_0), w(t_0+1), \ldots, w(T-1), w(T))$ are independent under $\mathbb{P}$ if $\mathbb{P}$ can be decompsed as a product

$$\mathbb{P} = \mu_{t_0} \otimes \cdots \otimes \mu_T$$

Equipping the set Ω of scenarios with a probability $\mathbb{P}$ is a delicate issue!

- The probabilistic distribution of the climate sensitivity parameter in climate models differs according to authors

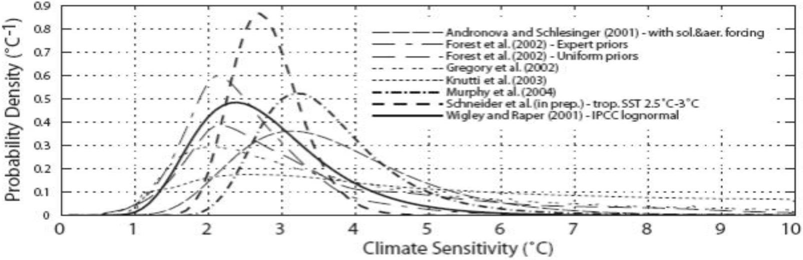

- In the multi-prior approach, the a priori information consists of different probabilities (*beliefs*, *priors*), belonging to a set $\mathcal{P}$ of admissible probabilities on Ω

In the set-membership approach, only a subset of the set of scenarios is known

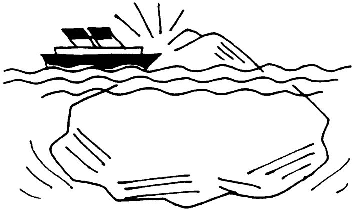

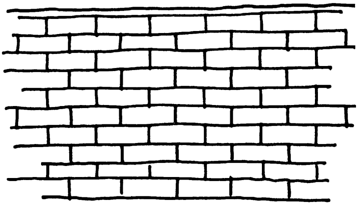

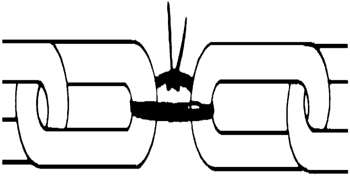

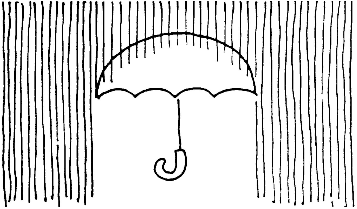

A priori information on the scenarios may be set membership

The general case

- Selected scenarios may belong to any subset $\overline{\Omega}$

$$w(\cdot) \in \overline{\Omega} \subset \Omega$$

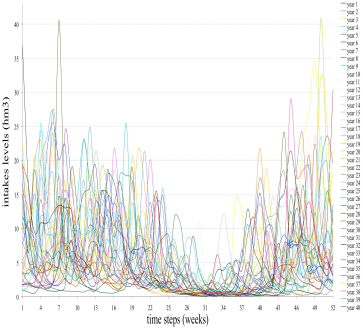

Historical water inflows scenarios in a dam

We can represent off-line information by the observed historical water inflows scenarios

Specific subsets correspond to time independence

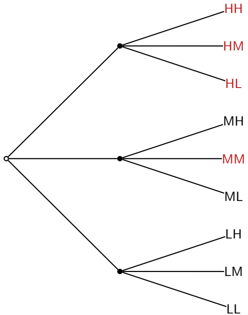

There is no time independence because the range of values of $w(t+1)$ depends on the value of $w(t)$:
$w(t) = H \Rightarrow w(t+1) \in \{M, L\}$
$w(t) = M \Rightarrow w(t+1) \in \{M\}$

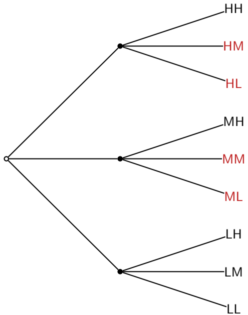

There is time independence because
$\overline{\Omega} = \{H, M\} \times \{M, L\} \subset \Omega$
is a product set

A priori information on the scenarios may be set membership

The product case

- Uncertain variables may be restricted to subsets, period by period

$$w(t) \in \mathbb{S}(t)$$

so that some scenarios are selected and the rest are excluded

$$w(\cdot) \in \mathbb{S}(t_0) \times \cdots \times \mathbb{S}(T) \subset \Omega = \mathbb{W}^{T+1-t_0}$$

Bounded water inflows in a dam

If only an upper bound on water inflows is known, we represent off-line information by

$$0 \leq a(t) \leq a^{\sharp}$$

A priori information on the scenarios may be softer than set membership thanks to plausibility functions

- The counterpart of a probability $\mathbb{P}$ — that weighs the likelihood of an event — is a plausibility function $\mathbb{Q}$
- Plausibility function $\mathbb{Q} : \Omega \to \mathbb{R} \cup \{-\infty\}$ can "soften" the above set membership approach
  - the higher $\mathbb{Q}\big(w(\cdot)\big)$, the more plausible the scenario $w(\cdot)$
  - totally implausible scenarios are those for which $\mathbb{Q}\big(w(\cdot)\big) = -\infty$

Historical water inflows scenarios in a dam

Attribute the value $\mathbb{Q}\big(w(\cdot)\big) = -\infty$ for all the scenarios $w(\cdot)$ which do not belong to the observed historical water inflows scenarios

The fear operator (Pierre Bernhard) is the robust counterpart of a probability

- Let $\mathbb{Q} : \Omega \to \mathbb{R} \cup \{-\infty\}$ be a a plausibility function
- The feared value of a function $A : \Omega \to \mathbb{R}$ is defined by

$$\mathbb{F}_{\mathbb{Q}}(A) := \min_{w(\cdot)\in\Omega} \left[A\big(w(\cdot)\big) - \mathbb{Q}\big(w(\cdot)\big) \right]$$

- The fear operator $\mathbb{F}_{\mathbb{Q}}$ enjoys linearity in the $(\min, +)$ algebra:

$$\mathbb{F}_{\mathbb{Q}}(\min\{A, B\}) = \min\{\mathbb{F}_{\mathbb{Q}}(A), \mathbb{F}_{\mathbb{Q}}(B)\}$$

- In the $(\min, +)$ algebra, the plausibility function $\mathbb{Q}$ plays the role of a weight, paralleling a probability distribution
- The uncertainties $(w(t_0), w(t_0+1), \ldots, w(T-1), w(T))$ are independent under $\mathbb{Q}$ if $\mathbb{Q}$ can be decomposed as a sum

$$\mathbb{Q} = \nu_{t_0} + \cdots + \nu_T$$

Summary

- A priori information is carried by the scenarios set, and may be
  - probabilistic
  - set membership
- This will be useful to mathematically express the objectives and the constraints in a decision problem under uncertainty

Outline of the presentation

Outline of the presentation

1. Dynamical control systems under uncertainty
   - Examples of uncertainties in dynamical systems
   - Uncertainty variables are new input variables
2. Scenarios support a priori/off-line information
   - Scenarios are temporal sequence of uncertainties
   - A priori / off-line information
3. On-line information feeds policies
   - The concept of policy
   - State and control solution maps
4. Summary

The issue of on-line information

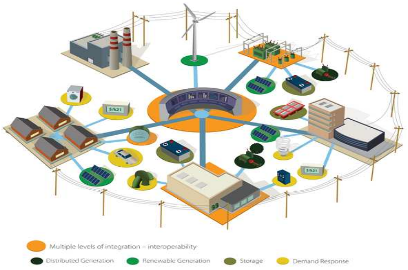

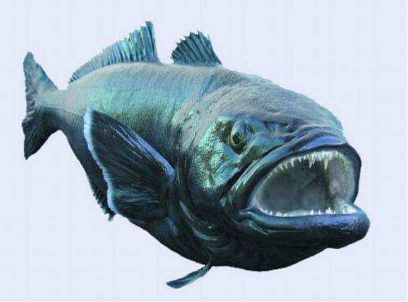

- Can we centralize all the informations on stock values in a large power system?
- Can we measure on-line the abundances of all age-classes in a population model?
- What about measurement errors?

When decisions do not take into account on-line information and the clock time, we are in the stationary static case

Stationary (open-loop)

Stationary open-loop control is

$$u : \underbrace{t \in \mathbb{T}}_{\text{time}} \mapsto \underbrace{u(t) \equiv u_{\mathrm{E}} \in \mathbb{U}}_{\text{control}}$$

Harvest the same biomass every year, as in the maximum sustainable yield

When decisions do not take into account on-line information but depend on the clock, we are in the open-loop case

Open-loop ⏲

Open-loop control consists of time-dependent sequences (planning, scheduling)

$$u : \underbrace{t \in \mathbb{T}}_{\text{time}} \mapsto \underbrace{u(t) \in \mathbb{U}}_{\text{control}}$$

Examples of open-loop control

- Fixed cycle gears for traffic lights in traffic regulation
- Mine planning: extract a given sequence of blocks every year, whatever you learn of the metal prices or of the ore content
- Solutions to optimal control problems by Pontryagin's variational approach

"I started work on control theory" (Richard Bellman)

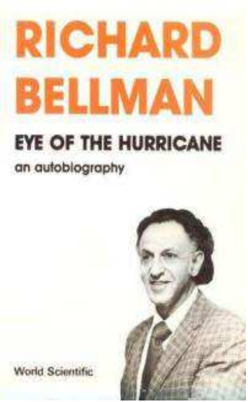

The tool we used was the calculus of variations. What we found was that very simple problems required great ingenuity. A small change in the problem caused a great change in the solution.

Clearly, something was wrong. There was an obvious lack of balance. Reluctantly, I was forced to the conclusion that the calculus of variations was not an effective tool for obtaining a solution

"The thought was finally forced upon me that the desired solution in a control process was a policy" (Richard Bellman)

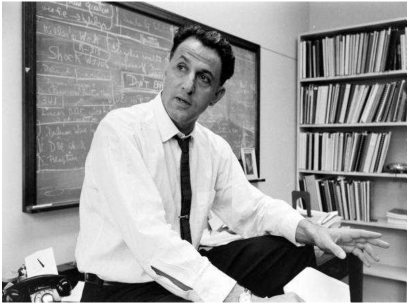

Richard Ernest Bellman
(August 26, 1920 – March 19, 1984)

From planning
to contingent planning

> *Again the intriguing thought: A solution is not merely a set of functions of time, or a set of numbers, but a rule telling the decisionmaker what to do; a policy (Richard Bellman)*

A computer code aboard a launcher embodies the concept of policy

```
if     state==0,
       do control=8
elseif state==1,
       do control=5.4
else do control=-15
```

On 4 June 1996, the maiden flight of the Ariane 5 launcher ended in a failure. (...) The attitude of the launcher and its movements in space are measured by an Inertial Reference System (SRI). (...) The data from the SRI are transmitted through the databus to the On-Board Computer (OBC), which executes the flight program (...) The Operand Error occurred due to an unexpected high value of an internal alignment function result called BH, Horizontal Bias, related to the horizontal velocity sensed by the platform

“The blind cat does not catch mice”

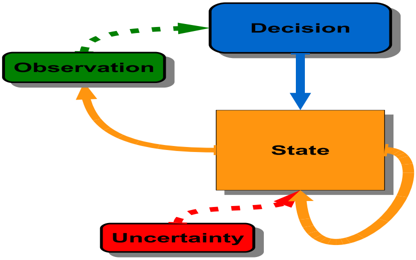

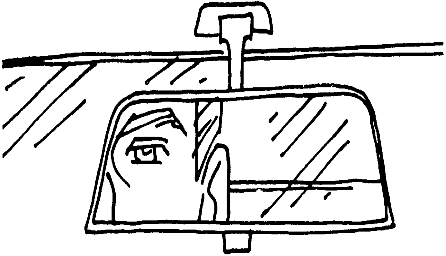

$$\underbrace{u(t)}_{\text{control}} = \texttt{Pol}\big(t, \underbrace{y(t)}_{\text{output}}\big)$$

- adaptive
- adjustable
- feeback
- wait and see
- full recourse
- on-line management
- corrective (vs. preventive)

How clouded the crystal ball looks beforehand

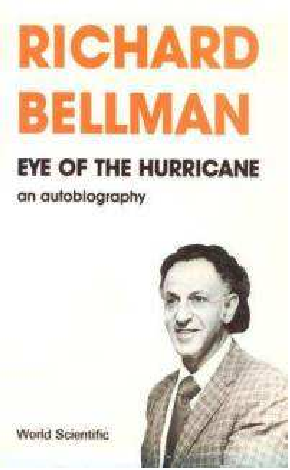

What is worth noting about the foregoing development is that I should have seen the application of dynamic programming to control theory several years before. I should have, but I didn't. It is very well to start a lecture by saying, 'Clearly, a control process can be regarded as a multistage decision process in which. . . ,' but it is a bit misleading.

Scientific developments can always be made logical and rational with sufficient hindsight. It is amazing, however, how clouded the crystal ball looks beforehand. We all wear such intellectual blinders and make such inexplicable blunders that it is amazing that any progress is made at all.

There are different observation patterns

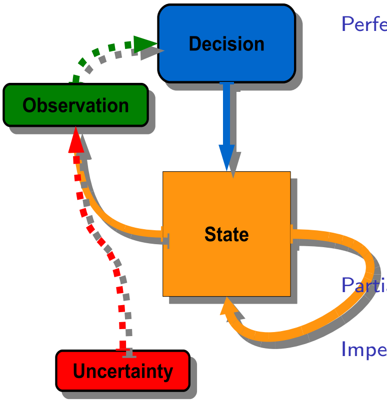

Perfect observation:

- Decision-hazard

$$y(t) = x(t)$$

- Hazard-decision

$$y(t) = \big(x(t), w(t)\big)$$

Partial observation:

$$y(t) = \texttt{Obs}\big(t, x(t)\big)$$

Imperfect observation:

$$y(t) = \texttt{Obs}\big(t, x(t), w(t)\big)$$

Dams management

Observing the stocks of all dams / the stocks and the water inflows / or only some stocks

State feedback policies correspond to perfect observation of the state

'Do thus-and-thus if you find yourself in this portion of state space with this amount of time left' (Richard Bellman)

Closed-loop control, state feedback (decision rule)

$$\texttt{Pol} : \underbrace{(t, x) \in \mathbb{T} \times \mathbb{X}}_{\text{(time, state)}} \mapsto u = \underbrace{\texttt{Pol}(t, x) \in \mathbb{U}}_{\text{control}}$$

Turbinate a fraction of the dam stock

$$\texttt{Pol}(t, S) = \alpha(t)S \ \text{ with } \ 0 \leq \alpha(t) \leq 1$$

ICES precautionary approach

$$\lambda_{UA}(N) = \max\{\lambda \in \mathbb{R}_+ \mid \texttt{SSB}(\texttt{Dyn}(N, \lambda)) \geq B_{\text{lim}} \ \text{ and } \ F(\lambda) \leq F_{\text{lim}}\}$$

Going from planning to contingent planning, we have considerably enlarged the set of solutions

- Stationary (open-loop): stationary sequences

$$u : t \in \mathbb{T} \mapsto u(t) \equiv u_{\mathrm{E}} \; , \quad u_{\mathrm{E}} \in \mathbb{U}$$

  Once the control space $\mathbb{U}$ is discretized in $N_{\mathbb{U}}$ elements,
  the solution space cardinality is $N_{\mathbb{U}}$

- Open-loop: time-dependent sequences (planning, scheduling)

$$u : t \in \mathbb{T} \mapsto u(t) \; , \quad u(\cdot) := \big(u(t_0), \ldots, u(T-1)\big) \in \mathbb{U}^{\mathbb{T}}$$

  With $N_{\mathbb{T}}$ time periods, the solution space cardinality is $N_{\mathbb{U}}^{N_{\mathbb{T}}}$

- Closed-loop: time and state-dependent sequences

$$\mathtt{Pol} : (t,x) \in \mathbb{T} \times \mathbb{X} \mapsto u = \mathtt{Pol}(t,x) \in \mathbb{U} \; , \quad \mathtt{Pol} \in \mathbb{U}^{\mathbb{T} \times \mathbb{X}}$$

  Once the state space $\mathbb{X}$ is discretized in $N_{\mathbb{X}}$ elements,
  the solution space cardinality is $N_{\mathbb{U}}^{N_{\mathbb{T}} \times N_{\mathbb{X}}}$

Admissible state feedback policies express control constraints

The control constraints case restricts policies to admissible policies

$$\mathcal{U}^{ad} = \{\texttt{Pol} \mid \texttt{Pol}(t,x) \in \mathbb{B}(t,x) \;, \quad \forall (t,x) \in \mathbb{T} \times \mathbb{X}\}$$

Dam management physical volume constraint

In a water reservoir, the output flow (control) cannot be more than the stock volume (state) and than a capacity constraint

$$0 \leq q(t) \leq S(t) \text{ and } 0 \leq q(t) \leq q^\sharp$$

Hence, an dam management policy of the form

$$\texttt{Pol}(t,S) = \max\{q^\flat, \min\{ \text{ function of } (t,S), q^\sharp, S\}\}$$

is admissible, where $0 \leq q^\flat \leq q^\sharp$ captures a requirement of minimal outflow (for biodiversity preservation in downward rivers, for instance)

Outline of the presentation

1. Dynamical control systems under uncertainty
   - Examples of uncertainties in dynamical systems
   - Uncertainty variables are new input variables
2. Scenarios support a priori/off-line information
   - Scenarios are temporal sequence of uncertainties
   - A priori / off-line information
3. On-line information feeds policies
   - The concept of policy
   - State and control solution maps
4. Summary

Along a given scenario, the system is deterministic

Une intelligence qui, à un instant donné, connaîtrait toutes les forces dont la nature est animée, la position respective des êtres qui la composent, si d'ailleurs elle était assez vaste pour soumettre ces données à l'analyse, embrasserait dans la même formule les mouvements des plus grands corps de l'univers, et ceux du plus léger atome. Rien ne serait incertain pour elle, et l'avenir comme le passé seraient présents à ses yeux.

Pierre-Simon Laplace,
Essai philosophique sur les probabilités

State and control solution maps are defined inductively along each scenario

Pick up

- a scenario $w(\cdot) = \big(w(t_0), w(t_0+1), \dots, w(T)\big) \in \Omega$
- a policy $\texttt{Pol} : \underbrace{(t,x) \in \mathbb{T} \times \mathbb{X}}_{\text{(time, state)}} \mapsto u = \underbrace{\texttt{Pol}(t,x) \in \mathbb{U}}_{\text{control}}$
- an initial state $x(t_0) = x_0 \in \mathbb{X}$

1. Plug the state $x(t_0)$ into the "machine" $\texttt{Pol}$ → initial decision $u(t_0) = \texttt{Pol}(t_0, x(t_0))$
2. Run the dynamics → second state $x(t_0+1) = \texttt{Dyn}(t_0, x(t_0), u(t_0), w(t_0))$
3. Second decision $u(t_0+1) = \texttt{Pol}(t_0+1, x(t_0+1))$
4. And so on $x(t_0+2) = \texttt{Dyn}(t_0+1, x(t_0+1), u(t_0+1), w(t_0+1))$
5. ...

State and control solution maps

Let be given

- a policy $\texttt{Pol} : \mathbb{T} \times \mathbb{X} \to \mathbb{U}$
- a scenario $w(\cdot) \in \Omega$
- and an initial state x_0 at initial time t_0

State solution map

The state solution map $X_{\texttt{Dyn}}[t_0, x_0, \texttt{Pol}, w(\cdot)]$ is the unique state path $x(\cdot) = (x(t_0), x(t_0+1), \dots, x(T))$ solution of dynamic

$$x(t+1) = \texttt{Dyn}\big(t, x(t), \texttt{Pol}\big(t, x(t)\big), w(t)\big) , \quad t = t_0, \dots, T-1$$

starting from the initial condition $x(t_0) = x_0$ at time t_0
and associated with policy $\texttt{Pol}$ and scenario $w(\cdot)$

The control solution map $U_{\texttt{Dyn}}[t_0, x_0, \texttt{Pol}, w(\cdot)]$ is the associated decision path $u(\cdot) = (u(t_0), u(t_0+1), \dots, u(T-1))$ where $u(t) = \texttt{Pol}\big(t, x(t)\big)$
Everything above extends to the hazard-decision case

What have we covered so far?

A policy is a decision rule

$$x(t+1) = \texttt{Dyn}\Big(t, x(t), \overbrace{u(t)}^{\texttt{Pol}(t,x(t))}, \underbrace{w(t)}_{\text{uncertainty}}\Big)$$

- On-line information feeds decisions through a policy, a strategy, a decision rule
- State feedback policies are natural solutions given by dynamic programming methods
- Once a policy is fixed, what is the fate of the state?
- This fate depends on the sequence of uncertainties crossed by the state

Outline of the presentation

Summary

- Uncertain variables $w(t)$ upon which the decision-maker (DM) has no control whatsoever
- Control variables $u(t)$ the values of which the DM can fix at any time within given sets
- Output variables: state, observation, etc.
  - State variables $x(t)$ are recursively constructed from a dynamics $\texttt{Dyn}$ by $x(t+1) = \texttt{Dyn}\big(t, x(t), u(t), w(t)\big)$
  - Observation variables $y(t) = \texttt{Obs}\big(t, x(t), w(t)\big)$
- On-line information feeds decisions through policies $\texttt{Pol} : (t, x) \mapsto \texttt{Pol}(t, x)$, giving $u(t) = \texttt{Pol}(t, x(t))$
- A scenario $w(\cdot)$ is a temporal sequence of uncertainties
- The set Ω of scenarios is endowed with off-line information
- How do we express the objectives and the constraints in a decision problem under uncertainty? How do we compare policies?