

Dynamic Programming

Michel DE LARA
CERMICS, École des Ponts ParisTech
France

École des Ponts ParisTech

November 2, 2020

Outline of the presentation

Multistage Stochastic Optimization

Dynamic Programming Without State

Dynamic Programming With State

Dynamic Programming With State and White Noise

Dynamic Programming With State and White Noise (Complements)

Outline of the presentation

Multistage Stochastic Optimization

Dynamic Programming Without State

Dynamic Programming With State

Dynamic Programming With State and White Noise

Dynamic Programming With State and White Noise (Complements)

Basic data

- ▶ Let $(\Omega, \mathcal{A}_\infty, \mathbb{P})$ be a probability space
- ▶ Let $T \in \mathbb{N}^*$ be the **horizon**
- ▶ For **stages** $t = 0, \dots, T - 1$, let \mathbb{U}_t be the **control set**, a measurable set equipped with σ -field \mathcal{U}_t
- ▶ For **stages** $t = 0, \dots, T$, let \mathbb{W}_t be the **uncertainty set**, a measurable set equipped with σ -fields \mathcal{W}_t

History space

For $t = 0, \dots, T$, we define

- ▶ the **history space** \mathbb{H}_t

$$\mathbb{H}_t = \mathcal{W}_0 \times \prod_{s=0}^{t-1} (\mathcal{U}_s \times \mathcal{W}_{s+1})$$

equipped with the **history field** \mathcal{H}_t

$$\mathcal{H}_t = \mathcal{W}_0 \otimes \bigotimes_{s=0}^{t-1} (\mathcal{U}_s \otimes \mathcal{W}_{s+1})$$

- ▶ A generic element $h_t \in \mathbb{H}_t$ is called a **history**

$$h_t = (w_0, u_0, w_1, u_1, w_2, \dots, u_{t-2}, w_{t-1}, u_{t-1}, w_t)$$

$$[h_t] = [(w_0, (u_s, w_{s+1})_{s=0, \dots, t-1})] = (w_0, \dots, w_t) = w_{[0:t]}$$

$$h_{s:t} = (u_r, w_{r+1})_{r=s-1, \dots, t-1} = (u_{s-1}, w_s, \dots, u_{t-1}, w_t)$$

$$[h_{s:t}] = [(u_r, w_{r+1})_{r=s-1, \dots, t-1}] = (w_s, \dots, w_t) = w_{[s:t]}$$

Noise process, noise filtration and adapted processes

- ▶ For $t = 0, \dots, T$, let

$$\mathbf{W}_t : \Omega \rightarrow \mathbb{W}_t$$

be a **random variable** taking values in \mathbb{W}_t (**noise**)

- ▶ We introduce the **past noises**, or noise process up to stage t as

$$\mathbf{W}_{[0:t]} = (\mathbf{W}_0, \dots, \mathbf{W}_t) \in \mathbb{W}_{[0:t]} = \prod_{s=0}^t \mathbb{W}_s$$

- ▶ We introduce the **filtration** $\mathcal{A} = (\mathcal{A}_t)_{t=0, \dots, T}$ defined by

$$\mathcal{A}_t = \sigma(\mathbf{W}_0, \dots, \mathbf{W}_t), \quad t = 0, \dots, T$$

- ▶ Let $\mathbb{L}^0(\Omega, \mathcal{A}, \prod_{s=0}^{T-1} \mathbb{U}_s)$ be the space of **\mathcal{A} -adapted processes** $(\mathbf{U}_0, \dots, \mathbf{U}_{T-1})$ with values in $\prod_{s=0}^{T-1} \mathbb{U}_s$, that is, such that

$$\sigma(\mathbf{U}_0) \subset \mathcal{A}_0, \dots, \sigma(\mathbf{U}_{T-1}) \subset \mathcal{A}_{T-1}$$

Multistage stochastic optimization problem

- ▶ Let be given a (cost) function

$$j : \mathbb{H}_T \rightarrow] - \infty, +\infty]$$

bounded below, and measurable with respect to the field \mathcal{H}_T

- ▶ We consider the multistage stochastic optimization problem

$$\min_{(\mathbf{U}_0, \dots, \mathbf{U}_{T-1}) \in \mathbb{L}_{\mathcal{A}}^0(\Omega, \prod_{s=0}^{T-1} \mathbb{U}_s)} \mathbb{E}[j(\mathbf{W}_0, \mathbf{U}_0, \mathbf{W}_1, \dots, \mathbf{U}_{T-1}, \mathbf{W}_T)]$$

- ▶ that is, with criterion $J : \mathbb{L}^0(\Omega, \mathcal{A}, \prod_{s=0}^{T-1} \mathbb{U}_s) \rightarrow] - \infty, +\infty]$ given by

$$J(\mathbf{U}_0, \dots, \mathbf{U}_{T-1}) = \mathbb{E}[j(\mathbf{W}_0, \mathbf{U}_0, \mathbf{W}_1, \dots, \mathbf{U}_{T-1}, \mathbf{W}_T)]$$

Doob Theorem

If every control set \mathbb{U}_t is a separable complete metric space, for $t = 0, \dots, T - 1$, the condition

$$(\mathbf{U}_0, \dots, \mathbf{U}_{T-1}) \in \mathbb{L}_{\mathcal{A}}^0(\Omega, \prod_{s=0}^{T-1} \mathbb{U}_s)$$

is equivalent to the existence of **measurable mappings**

$$\lambda_t : \mathbb{W}_{0:t} \rightarrow \mathbb{U}_t, \quad t = 0, \dots, T - 1$$

such that

$$\mathbf{U}_t = \lambda_t(\mathbf{W}_{0:t}) = \lambda_t(\mathbf{W}_0, \dots, \mathbf{W}_t), \quad t = 0, \dots, T - 1$$

Curse of scenario expansion

Assuming that

- ▶ the control u_t can take N_u values
- ▶ the uncertainty w_t can take N_w values

we obtain

- ▶ N_w^T scenarios
- ▶ $1 + N_w + \dots + N_w^T$ nodes in the scenario tree
- ▶ $N_u \times \frac{N_w^{T+1} - 1}{N_w - 1} \approx N_u N_w^T$ elements in the solution space

so that the **number of possible solutions**
grows **exponentially** with the **number T of stages**

Complexity of upper and lower bounds

- ▶ Upper bound: open loop solution
- ▶ Lower bound: anticipative

Outline of the presentation

Multistage Stochastic Optimization

Dynamic Programming Without State

Dynamic Programming With State

Dynamic Programming With State and White Noise

Dynamic Programming With State and White Noise (Complements)

Bellman operators

For $t = 0, \dots, T$,

- ▶ we define $\mathbb{L}^\infty(\mathbb{H}_t, \mathcal{H}_t)$, the space of bounded measurable real-valued functions $\mathbb{H}_t \rightarrow \mathbb{R}$
- ▶ We suppose that there exists a regular conditional distribution

$$\mathbb{P}_{\mathbf{W}_{t+1}}^{\mathbf{W}_{[0:t]}}(w_{[0:t]}, dw_{t+1}) = \mathbb{P}_{\mathbf{W}_{t+1}}^{\mathbf{W}_{[0:t]}}([h_t], dw_{t+1})$$

of the random variable \mathbf{W}_{t+1} knowing the random process $\mathbf{W}_{[0:t]}$

- ▶ we define the **Bellman operators**

$\mathcal{B}_{t+1} : \mathbb{L}^\infty(\mathbb{H}_{t+1}, \mathcal{H}_{t+1}) \rightarrow \mathbb{L}^\infty(\mathbb{H}_t, \mathcal{H}_t)$ by

$$(\mathcal{B}_{t+1}\varphi)(h_t) = \inf_{u_t \in \mathbb{U}_t} \int_{\mathbb{W}_{t+1}} \varphi((h_t, u_t, w_{t+1})) \mathbb{P}_{\mathbf{W}_{t+1}}^{\mathbf{W}_{[0:t]}}([h_t], dw_{t+1})$$

$$\forall \varphi \in \mathbb{L}^\infty(\mathbb{H}_{t+1}, \mathcal{H}_{t+1}), \quad \forall h_t \in \mathbb{H}_t$$

Value functions and Bellman equation

- ▶ We define inductively **value functions**, or **Bellman functions**,

$$V_t : \mathbb{H}_t \rightarrow \mathbb{R}, \quad t = 0, \dots, T$$

by

$$V_T = j, \quad V_t = \mathcal{B}_{t+1} V_{t+1}, \quad t = 0, \dots, T - 1$$

- ▶ that is, solution of the **Bellman equation**

$$V_t(h_t) = \inf_{u_t \in \mathbb{U}_t} \int_{\mathbb{W}_{t+1}} V_{t+1}(h_t, u_t, w_{t+1}) \mathbb{P}_{\mathbb{W}_{t+1}}^{\mathbb{W}_{[0:t]}}([h_t], dw_{t+1})$$

Measurable selection

We suppose that, for all $t = 0, \dots, T$, there exists a **measurable selection**

$$\gamma_t : (\mathbb{H}_t, \mathcal{H}_t) \rightarrow (\mathbb{U}_t, \mathcal{U}_t)$$

such that

$$\gamma_t(h_t) \in \arg \min_{u_t \in \mathbb{U}_t} \int_{\mathbb{W}_{t+1}} V_{t+1}(h_t, u_t, w_{t+1}) \mathbb{P}_{\mathbb{W}_{t+1}}^{\mathbf{W}^{[0:t]}}([h_t], dw_{t+1})$$

$$\forall h_t \in \mathbb{H}_t$$

Proposition

A solution to the *multistage stochastic optimization problem*

$$\min_{(\mathbf{U}_0, \dots, \mathbf{U}_{T-1}) \in \mathbb{L}_{\mathcal{A}}^0(\Omega, \prod_{s=0}^{T-1} \mathbb{U}_s)} \mathbb{E}[j(\mathbf{W}_0, \mathbf{U}_0, \mathbf{W}_1, \dots, \mathbf{U}_{T-1}, \mathbf{W}_T)]$$

is the sequence $\mathbf{U}_0^*, \dots, \mathbf{U}_{T-1}^*$ of random variables defined inductively by

$$\mathbf{U}_t^* = \gamma_t \circ \mathbf{H}_t^*, \quad t = 0, \dots, T-1$$

$$\text{where } \mathbf{H}_0^* = \mathbf{W}_0, \quad \mathbf{H}_{t+1}^* = (\mathbf{H}_t^*, \mathbf{U}_t^*, \mathbf{W}_{t+1}), \quad t = 0, \dots, T-1$$

and the *minimum* is

$$\mathbb{E}[V_0(\mathbf{W}_0)] = \min_{(\mathbf{U}_0, \dots, \mathbf{U}_{T-1}) \in \mathbb{L}_{\mathcal{A}}^0(\Omega, \prod_{s=0}^{T-1} \mathbb{U}_s)} \mathbb{E}[j(\mathbf{W}_0, \mathbf{U}_0, \mathbf{W}_1, \dots, \mathbf{U}_{T-1}, \mathbf{W}_T)]$$

Extension

$$\mathbf{H}_0 = \mathbf{W}_0, \quad \mathbf{H}_{t+1} = (\mathbf{H}_t, \mathbf{U}_t, \mathbf{W}_{t+1})$$

Constraints of the form

$$(\mathbf{H}_t, \mathbf{U}_t) \in \mathbb{C}_t \subset \mathbb{H}_t \times \mathbb{U}_t, \quad \mathbb{P} - \text{a.s.}, \quad t = 0, \dots, T - 1$$

and

$$\mathbf{H}_T \in \mathbb{C}_T \subset \mathbb{H}_T, \quad \mathbb{P} - \text{a.s.}$$

Outline of the presentation

Multistage Stochastic Optimization


Dynamic Programming Without State

Dynamic Programming With State

Dynamic Programming With State and White Noise

Dynamic Programming With State and White Noise (Complements)

Optimal single dam management



A single dam nonlinear dynamical model in decision-hazard

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

$$\underbrace{S_{t+1}}_{\text{future volume}} = \min\{S^\#, \underbrace{S_t}_{\text{volume}} - \underbrace{Q_t}_{\text{turbined}} + \underbrace{A_{t+1}}_{\text{inflow volume}}\}$$

- ▶ S_t volume (stock) of water at the beginning of period $[t, t + 1[$
- ▶ A_{t+1} inflow water volume (rain, etc.) during $[t, t + 1[$
- ▶ decision-hazard:
 - A_{t+1} is not available at the beginning of period $[t, t + 1[$
- ▶ Q_t turbined outflow volume during $[t, t + 1[$
 - ▶ decided at the beginning of period $[t, t + 1[$
 - ▶ supposed to depend on S_t but not on A_{t+1}
 - ▶ chosen such that $0 \leq Q_t \leq S_t$

The traditional economic problem is maximizing the expected payoff

- ▶ Suppose that
 - ▶ a **probability** \mathbb{P} is given on the set \mathbb{R}^T of **water inflows scenarios** (A_0, \dots, A_{T-1})
 - ▶ **turbined water** Q_t is sold at **price** p_t , related to the price at which energy can be sold at time t
 - ▶ at the horizon, the **final volume** S_T has a **value** $K(S_T)$, the “final value of water”
- ▶ The traditional economic problem is to maximize the intertemporal payoff (without discounting if the horizon is short)

$$\max \mathbb{E} \left[\sum_{t=0}^{T-1} \underbrace{p_t Q_t}_{\text{turbined water payoff}} + \underbrace{K(S_T)}_{\text{final volume utility}} \right]$$

State reduction and dynamics

For $t = 0, \dots, T$, suppose that there exists

- ▶ **state space** \mathbb{X}_t , a measurable set equipped with σ -field \mathcal{X}_t
- ▶ **reduction mappings**

$$\theta_t : \mathbb{H}_t \rightarrow \mathbb{X}_t$$

- ▶ **dynamics**

$$f_t : \mathbb{X}_t \times \mathbb{U}_t \times \mathbb{W}_{t+1} \rightarrow \mathbb{X}_{t+1}$$

such that

$$\theta_{t+1}(h_t, u_t, w_{t+1}) = f_t(\theta_t(h_t), u_t, w_{t+1}), \quad t = 0, \dots, T - 1$$

Cost only depends on final state

Suppose that there exists

$$\tilde{j} : \mathbb{X}_T \rightarrow]-\infty, +\infty]$$

such that the cost function $j : \mathbb{H}_T \rightarrow]-\infty, +\infty]$ can be factored as

$$j = \tilde{j} \circ \theta_T$$

Markovian assumption

- ▶ Let $\Delta(\mathbb{W}_t)$ denote the set of probabilities on $(\mathbb{W}_t, \mathbb{W}_t)$, for $t = 0, \dots, T$
- ▶ Suppose that, for all $t = 0, \dots, T$, there exists

$$\mu_t : \mathbb{X}_t \times \prod_{s=0}^t \mathbb{W}_s \rightarrow \Delta(\mathbb{W}_{t+1})$$

such that

$$\mathbb{P}_{\mathbf{w}_{t+1}}^{\mathbf{W}^{[0:t]}}([h_t], dw_{t+1}) = \mu_t(\theta_t(h_t), dw_{t+1})$$

Bellman equation

- ▶ We define inductively

$$\tilde{V}_T(x_T) = \tilde{j}(x_T), \quad \forall x_T \in \mathbb{X}_T$$

$$\tilde{V}_t(x_t) = \inf_{u_t \in \mathbb{U}_t} \int_{\mathbb{W}_{t+1}} \tilde{V}_{t+1}(f_t(x_t, u_t, w_{t+1})) \mu_t(x_t, dw_{t+1})$$

$$\forall x_t \in \mathbb{X}_t, \quad t = 0, \dots, T-1$$

- ▶ We suppose that there exists a measurable selection

$$\tilde{\gamma}_t^* : (\mathbb{X}_t, \mathcal{X}_t) \rightarrow (\mathbb{U}_t, \mathcal{U}_t), \quad t = 0, \dots, T-1$$

such that

$$\tilde{\gamma}_t^*(x_t) \in \arg \min_{u_t \in \mathbb{U}_t} \int_{\mathbb{W}_{t+1}} \tilde{V}_{t+1}(f_t(x_t, u_t, w_{t+1})) \mu_t(x_t, dw_{t+1})$$

$$\forall x_t \in \mathbb{X}_t$$

Proposition

A solution to the *multistage stochastic optimization problem*

$$\min_{\mathbf{U}_0, \dots, \mathbf{U}_{T-1}} \mathbb{E}[j(\mathbf{W}_0, \mathbf{U}_0, \mathbf{W}_1, \dots, \mathbf{U}_{T-1}, \mathbf{W}_T)]$$
$$\sigma(\mathbf{U}_0) \subset \sigma(\mathbf{W}_0), \dots, \sigma(\mathbf{U}_{T-1}) \subset \sigma(\mathbf{W}_0, \dots, \mathbf{W}_{T-1})$$

is the sequence $\mathbf{U}_0^*, \dots, \mathbf{U}_{T-1}^*$ of random variables defined inductively by

$$\mathbf{U}_t^* = \tilde{\gamma}_t^*(\mathbf{X}_t^*), \quad t = 0, \dots, T-1$$

$$\text{where } \mathbf{X}_0^* = \mathbf{W}_0, \quad \mathbf{X}_{t+1}^* = f_t(\mathbf{X}_t^*, \mathbf{U}_t^*, \mathbf{W}_{t+1}), \quad t = 0, \dots, T-1$$

and the *minimum* is

$$\mathbb{E}[\tilde{V}_0(\mathbf{X}_0^*)] = \min_{(\mathbf{U}_0, \dots, \mathbf{U}_{T-1}) \in \mathbb{L}_{\mathcal{A}}^0(\Omega, \prod_{s=0}^{T-1} \mathbb{U}_s)} \mathbb{E}[j(\mathbf{W}_0, \mathbf{U}_0, \mathbf{W}_1, \dots, \mathbf{U}_{T-1}, \mathbf{W}_T)]$$

Extension

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

Constraints of the form

$$(\mathbf{X}_t, \mathbf{U}_t) \in \mathbb{C}_t \subset \mathbb{X}_t \times \mathbb{U}_t, \quad \mathbb{P} - \text{a.s.}, \quad t = 0, \dots, T-1$$

and

$$\mathbf{X}_T \in \mathbb{C}_T \subset \mathbb{X}_T, \quad \mathbb{P} - \text{a.s.}$$

Outline of the presentation

Multistage Stochastic Optimization

Dynamic Programming Without State

Dynamic Programming With State

Dynamic Programming With State and White Noise

Dynamic Programming With State and White Noise (Complements)

Stochastic optimal control problem formulation

$$\min_{\mathbf{u}_0, \dots, \mathbf{u}_{T-1}} \mathbb{E} \left[\sum_{t=0}^{T-1} L_t(\mathbf{x}_t, \mathbf{u}_t, \mathbf{w}_{t+1}) + K(\mathbf{x}_T) \right]$$

$$\sigma(\mathbf{u}_0) \subset \sigma(\mathbf{x}_0), \dots, \sigma(\mathbf{u}_{T-1}) \subset \sigma(\mathbf{x}_0, \mathbf{w}_1, \dots, \mathbf{w}_{T-1})$$

$$\mathbf{x}_{t+1} = f_t(\mathbf{x}_t, \mathbf{u}_t, \mathbf{w}_{t+1}), \quad t = 0, \dots, T-1$$

$$\mathbf{u}_t \in \mathbb{B}_t(\mathbf{x}_t), \quad t = 0, \dots, T-1$$

Basic data

Let \mathbb{U} , \mathbb{W} , \mathbb{X} be measurable sets, equipped with σ -fields \mathcal{U} , \mathcal{W} , \mathcal{X} and, for $t = 0, \dots, T - 1$,

- ▶ **dynamics** mapping

$$f_t : \mathbb{X} \times \mathbb{U} \times \mathbb{W} \rightarrow \mathbb{X}$$

- ▶ **instantaneous costs** functions

$$L_t : \mathbb{X} \times \mathbb{U} \times \mathbb{W} \rightarrow \mathbb{R}$$

- ▶ **final cost** function

$$K : \mathbb{X} \rightarrow \mathbb{R}$$

- ▶ **constraints** set-valued mapping

$$\mathbb{B}_t : \mathbb{X} \rightrightarrows \mathbb{U}$$

Bellman equation

- ▶ We consider a stochastic process $(\mathbf{W}_1, \dots, \mathbf{W}_T)$, with values in \mathbb{W}
- ▶ We define inductively the **Bellman functions**

$$V_T(x) = K(x), \quad \forall x \in \mathbb{X}$$

$$V_t(x) = \inf_{u \in \mathbb{B}_t(x)} \mathbb{E}_{\mathbf{W}_{t+1}} [L_t(x, u, \mathbf{W}_{t+1}) + V_{t+1}(f_t(x, u, \mathbf{W}_{t+1}))]$$

$$\forall x \in \mathbb{X}, \quad t = 0, \dots, T - 1$$

- ▶ We suppose that there exists a measurable selection

$$\gamma_t^* : (\mathbb{X}, \mathcal{X}) \rightarrow (\mathbb{U}, \mathcal{U}), \quad t = 0, \dots, T - 1 \text{ such that}$$

$$\gamma_t^*(x) \in \arg \min_{u \in \mathbb{B}_t(x)} \mathbb{E}_{\mathbf{W}_{t+1}} [L_t(x, u, \mathbf{W}_{t+1}) + V_{t+1}(f_t(x, u, \mathbf{W}_{t+1}))]$$

$$\forall x \in \mathbb{X}$$

White noise assumption

- ▶ We suppose that the stochastic process $(\mathbf{W}_1, \dots, \mathbf{W}_T)$ is a **white noise**, that is, $\mathbf{W}_1, \dots, \mathbf{W}_T$ are independent random variables
- ▶ We consider a **random variable** \mathbf{X}_0 , with values in \mathbb{X} , **independent of** the stochastic process $(\mathbf{W}_1, \dots, \mathbf{W}_T)$

Bellman result

Proposition

A solution to the *multistage stochastic optimization problem*

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

$$\sigma(\mathbf{U}_0) \subset \sigma(\mathbf{X}_0), \dots, \sigma(\mathbf{U}_{T-1}) \subset \sigma(\mathbf{X}_0, \mathbf{W}_1, \dots, \mathbf{W}_{T-1})$$

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

$$\mathbf{U}_t \in \mathbb{B}_t(\mathbf{X}_t), \quad t = 0, \dots, T-1$$

is the sequence $\mathbf{U}_0^*, \dots, \mathbf{U}_{T-1}^*$ of random variables defined inductively by

$$\mathbf{U}_t^* = \gamma_t^*(\mathbf{X}_t^*), \quad t = 0, \dots, T-1$$

where $\mathbf{X}_0^* = \mathbf{X}_0$, $\mathbf{X}_{t+1}^* = f_t(\mathbf{X}_t^*, \mathbf{U}_t^*, \mathbf{W}_{t+1})$, $t = 0, \dots, T-1$

and the *minimum* is

$$\mathbb{E} [V_0(\mathbf{X}_0)] = \min_{\mathbf{U}_0, \dots, \mathbf{U}_{T-1}} \mathbb{E} \left[\sum_{t=0}^{T-1} L_t(\mathbf{X}_t, \mathbf{U}_t, \mathbf{W}_{t+1}) + K(\mathbf{X}_T) \right]$$

Outline of the presentation

Multistage Stochastic Optimization

Dynamic Programming Without State

Dynamic Programming With State

Dynamic Programming With State and White Noise

Dynamic Programming With State and White Noise (Complements)

Outline of the presentation

Multistage Stochastic Optimization

Dynamic Programming Without State

Dynamic Programming With State

Dynamic Programming With State and White Noise

Dynamic Programming With State and White Noise (Complements)

Bellman's Principle of Optimality


Cost-to-go and Bellman equation

Backward offline / forward online

The curse of dimensionality

Hazard-decision, linear-convex, SDDP


The shortest path on a graph illustrates Bellman's Principle of Optimality



*For an auto travel analogy, suppose that the fastest route from **Los Angeles** to to **Boston** passes through **Chicago**.*

*The principle of optimality translates to obvious fact that the **Chicago to Boston** portion of the route is also the fastest route for a trip that starts from **Chicago** and ends in **Boston**.
(Dimitri P. Bertsekas)*

Bellman's Principle of Optimality



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

An optimal policy has the property that whatever the initial state and initial decision are, the remaining decisions must constitute an optimal policy with regard to the state resulting from the first decision (Richard Bellman)

What is state and what is noise?

Delineating what is state and what is noise is a modelling issue

When the uncertainties are not independent,
a solution is to enlarge the state

- ▶ If the water inflows follow an auto-regressive model, we have

$$\begin{array}{l} \text{future stock} \\ \underbrace{S_{t+1}} \\ \text{future water inflows} \\ \underbrace{A_{t+1}} \end{array} = \min \left\{ S^\#, \underbrace{S_t}_{\text{stock}} - \underbrace{Q_t}_{\text{water release}} + \underbrace{A_{t+1}}_{\text{water inflows}} \right\}$$
$$\underbrace{A_{t+1}}_{\text{future water inflows}} = \alpha \underbrace{A_t}_{\text{water inflows}} + \underbrace{W_{t+1}}_{\text{noise}}$$

where we suppose that $\mathbf{W}_1, \dots, \mathbf{W}_{T-1}, \mathbf{W}_T$
form a sequence of **independent** random variables

- ▶ The couple $x_t = (S_t, A_t)$ is a **sufficient summary** of past controls and uncertainties to do forecasting:
knowing the **state** $x_t = (S_t, A_t)$ at time t is **sufficient**
to forecast x_{t+1} , given the control Q_t and the uncertainty \mathbf{W}_{t+1}

What is a state?


Bellman autobiography, Eye of the Hurricane

Conversely, once it was realized that the concept of policy was fundamental in control theory, the mathematicization of the basic engineering concept of 'feedback control,' then the emphasis upon a state variable formulation became natural

- ▶ A state in optimal stochastic control problems is a sufficient statistics for the uncertainties and past controls (P. Whittle, *Optimization over Time: Dynamic Programming and Stochastic Control*)
- ▶ Quoting Whittle, suppose there is a variable x_t which summarizes past history in that, given t and the value of x_t , one can calculate the optimal u_t and also x_{t+1} without knowledge of the history $(\omega, u_0, \dots, u_{t-1})$, for all t , where ω represents all uncertainties. Such a variable is termed *sufficient*
- ▶ While history takes value in an increasing space as t increases, a sufficient variable taking values in a space independent of t is called a state variable


A bit of history (and fun)

“Where did the name, dynamic programming, come from?”



The 1950s were not good years for mathematical research. We had a very interesting gentleman in Washington named Wilson. He was Secretary of Defense, and he actually had a pathological fear and hatred of the word, research. I'm not using the term lightly; I'm using it precisely. His face would suffuse, he would turn red, and he would get violent if people used the term, research, in his presence. You can imagine how he felt, then, about the term, mathematical.


“Where did the name, dynamic programming, come from?”



What title, what name, could I choose? In the first place I was interested in planning, in decision making, in thinking. But planning, is not a good word for various reasons. I decided therefore to use the word, programming.

“Where did the name, dynamic programming, come from?”

I wanted to get across the idea that this was dynamic, this was multistage, this was time-varying. I thought, let's kill two birds with one stone. Let's take a word that has an absolutely precise meaning, namely dynamic, in the classical physical sense. It also has a very interesting property as an adjective, and that is it's impossible to use the word, dynamic, in a pejorative sense. Try thinking of some combination that will possibly give it a pejorative meaning. It's impossible. Thus, I thought dynamic programming was a good name.



Outline of the presentation

Multistage Stochastic Optimization

Dynamic Programming Without State

Dynamic Programming With State

Dynamic Programming With State and White Noise

Dynamic Programming With State and White Noise (Complements)

Bellman's Principle of Optimality

Cost-to-go and Bellman equation

Backward offline / forward online

The curse of dimensionality

Hazard-decision, linear-convex, SDDP

The cost-to-go / value function / Bellman function

Assume that the primitive random variables $\mathbf{W}_1, \dots, \mathbf{W}_{T-1}, \mathbf{W}_T$ are independent under the probability \mathbb{P}

Cost-to-go / value function / Bellman function

The cost-to-go from state x at stage t is

$$V_t(x) = \min_{\gamma_t, \dots, \gamma_{T-1}} \mathbb{E} \left[\sum_{s=t}^{T-1} L_s(\mathbf{X}_s, \mathbf{U}_s, \mathbf{W}_{s+1}) + K(\mathbf{X}_T) \right]$$

where $\mathbf{X}_t = x$ and, for $s = t, \dots, T-1$,
 $\mathbf{X}_{s+1} = f_s(\mathbf{X}_s, \mathbf{U}_s, \mathbf{W}_{s+1})$ and $\mathbf{U}_s = \gamma_s(\mathbf{X}_s)$

- ▶ The function $V_t : \mathbb{X} \rightarrow \overline{\mathbb{R}}$ is called the value function, or the Bellman function
- ▶ The original minimization problem is $V_0(x_0)$

The stochastic dynamic programming equation, or Bellman equation, is a backward equation satisfied by the value function

Stochastic dynamic programming equation

If the primitive random variables $\mathbf{W}_1, \dots, \mathbf{W}_{T-1}, \mathbf{W}_T$ are independent under the probability \mathbb{P} , the value function $V_t : \mathbb{X} \rightarrow \overline{\mathbb{R}}$ satisfies the following backward induction, where t runs from $T - 1$ down to 0

$$V_T(x) = K(x)$$


$$V_t(x) = \min_{u \in \mathbb{B}_t(x)} \mathbb{E}_{\mathbf{W}_{t+1}} \left[L_t(x, u, \mathbf{W}_{t+1}) + V_{t+1}(f_t(x, u, \mathbf{W}_{t+1})) \right]$$

$$\forall x \in \mathbb{X}$$

Sketch of the proof in the deterministic case

$$V_t(x) = \min_{u \in \mathbb{B}_t(x)} \left(\underbrace{L_t(x, u)}_{\text{instantaneous cost}} + \overbrace{V_{t+1}(f_t(x, u))}^{\text{optimal cost}} \right)$$

future state



A decision u at time t in state x provides

- ▶ an instantaneous cost $L_t(x, u)$
- ▶ and a future cost for attaining the new state $f_t(x, u)$

Outline of the presentation

Multistage Stochastic Optimization

Dynamic Programming Without State

Dynamic Programming With State

Dynamic Programming With State and White Noise

Dynamic Programming With State and White Noise (Complements)

Bellman's Principle of Optimality

Cost-to-go and Bellman equation

Backward offline / forward online

The curse of dimensionality

Hazard-decision, linear-convex, SDDP

Navigating between “backward offline” and “forward online”

Optimal trajectories are calculated forward online

1. Initial state $x_0^* = x_0$
2. Plug the state x_0^* into the feedback γ_0
→ initial decision $u_0^* = \gamma_0^*(x_0^*)$
or compute the **optimal decision** u_t^* “on the fly” by


$$u_0^* \in \arg \min_{u \in \mathbb{B}_0(x_0^*)} \mathbb{E}_{\mathbf{W}_1} \left[L_0(x_0^*, u, \mathbf{W}_1) + V_1(f_0(x_0^*, u, \mathbf{W}_1)) \right]$$

3. Run the dynamics → second state $x_1^* = f_0(x_0^*, u_0^*, w_1)$
4. Second decision

$$u_1^* \in \arg \min_{u \in \mathbb{B}_1(x_1^*)} \mathbb{E}_{\mathbf{W}_2} \left[L_1(x_1^*, u, \mathbf{W}_2) + V_2(f_1(x_1^*, u, \mathbf{W}_2)) \right]$$

5. And so on $x_2^* = f_1(x_1^*, u_1^*, w_2)$
6. ...

“Life is lived forward but understood backward”
(Søren Kierkegaard)



D. P. Bertsekas introduces his book
Dynamic Programming and Optimal Control
with a citation by Søren Kierkegaard

*“Livet skal forstås baglaens, men
laves forlaens”*

*Life is to be understood backwards,
but it is lived forwards*

- ▶ The value function and the optimal policies are computed **backward** and **offline** by means of the Bellman equation
- ▶ whereas the optimal trajectories are computed **forward** and **online**

How optimal decisions can be computed online

Greedy one-step lookahead algorithm

- ▶ If we are able to store the value functions $x \mapsto V_t(x)$
- ▶ we do not need to compute and store the optimal policy γ_t^* in advance
- ▶ Indeed, when we are at state x at time t in real time, we can just compute the **optimal decision** u_t^* “on the fly” by

$$u_t^* \in \arg \min_{u \in \mathbb{B}_t(x)} \mathbb{E}_{\mathbf{W}_{t+1}} \left[L_t(x, u, \mathbf{W}_{t+1}) + V_{t+1}(f_t(x, u, \mathbf{W}_{t+1})) \right]$$

- ▶ In addition to sparing storage, this method makes it possible to incorporate in the above program any new information available at time t (on the distribution of the noise \mathbf{W}_{t+1} , for instance)

So, the question is: how can we store the value functions?

The effort can be concentrated on computing the value functions

- ▶ on a grid, by discretizing the Bellman equation
(but curse of dimensionality)
- ▶ by estimating basis coefficients,
when it is known that the value functions are quadratic
(the linear-quadratic case)
- ▶ by estimating lower affine approximations of the value functions,
when it is known that the value function is convex
(the linear-convex case and SDDP)

Outline of the presentation

Multistage Stochastic Optimization

Dynamic Programming Without State

Dynamic Programming With State

Dynamic Programming With State and White Noise

Dynamic Programming With State and White Noise (Complements)

Bellman's Principle of Optimality

Cost-to-go and Bellman equation

Backward offline / forward online

The curse of dimensionality

Hazard-decision, linear-convex, SDDP

The curse of dimensionality :- (

Algorithm for the Bellman functions

```
initialization  $V_T = K$ ;  
for  $t = T, T - 1, \dots, 0$  do  
  forall  $x \in \mathbb{X}$  do  
    forall  $u \in \mathbb{B}_t(x)$  do  
      forall  $w \in \mathbb{W}_{t+1}$  do  
         $l_t(x, u, w) = L_t(x, u, w) + V_{t+1}(f_t(x, u, w))$   
         $\sum_{w \in \mathbb{W}_{t+1}} \mathbb{P}\{w\} l_t(x, u, w)$   
  
   $V_t(x) = \min_{u \in \mathbb{B}_t(x)} \sum_{w \in \mathbb{W}_{t+1}} \mathbb{P}\{w\} l_t(x, u, w) ;$   
  
   $\mathbb{B}_t^*(x) = \arg \min_{u \in \mathbb{B}_t(x)} \sum_{w \in \mathbb{W}_{t+1}} \mathbb{P}\{w\} l_t(x, u, w)$ 
```


Complexity of the dynamic programming algorithm

Assuming that

- ▶ the state x_t can take N_x values
- ▶ the control u_t can take N_u values
- ▶ the uncertainty w_t can take N_w values

the complexity (number of operations) of the Bellman algorithm is in

$$O(T \times N_x \times N_u \times N_w)$$

which is linear in the number of stages :-)

but exponential in the **dimension of the state**
(and also control and uncertainty)

The curse of dimensionality is illustrated by the random access memory capacity on a computer: one, two, three, infinity (Gamov)

- ▶ On a computer
 - ▶ RAM: 8 GBytes = $8(1\ 024)^3 = 2^{33}$ bytes
 - ▶ a double-precision real: 8 bytes = 2^3 bytes
 - ▶ $\implies 2^{30} \approx 10^9$ double-precision reals can be handled in RAM
- ▶ If a state of dimension 4 is approximated by a grid with 100 levels by components, we need to manipulate $100^4 = 10^8$ reals and
 - ▶ do a time loop
 - ▶ do a control loop (after discretization)
 - ▶ compute an expectation

The wall of dimension can be pushed beyond
if additional properties are exploited (linearity, convexity)

In the linear-quadratic case, value functions are quadratic and optimal policies are linear

- ▶ When cost functions are quadratic

$$\begin{aligned}K(x) &= x' S_T x (+\text{affine}) \\L_t(x, u, w) &= x' S_t x + w' R_t w + u' Q_t u (+\text{affine})\end{aligned}$$

- ▶ and the dynamic is affine

$$f_t(x, u, w) = F_t x + G_t u + H_t w (+\text{constant})$$

- ▶ and primitive random variables $\mathbf{W}_1, \dots, \mathbf{W}_{T-1}, \mathbf{W}_T$ are square integrable and independent under the probability \mathbb{P}
- ▶ then, the value functions $x \mapsto V_t(x)$ are quadratic, and optimal policies are affine in the state

$$u_t = M_t x_t (+\text{constant})$$

Outline of the presentation

Multistage Stochastic Optimization

Dynamic Programming Without State

Dynamic Programming With State

Dynamic Programming With State and White Noise

Dynamic Programming With State and White Noise (Complements)

Bellman's Principle of Optimality

Cost-to-go and Bellman equation

Backward offline / forward online

The curse of dimensionality

Hazard-decision, linear-convex, SDDP

Stochastic optimal control problem formulation (hazard-decision)

$$\min_{\mathbf{u}_0, \dots, \mathbf{u}_{T-1}} \mathbb{E} \left[\sum_{t=0}^{T-1} L_t(\mathbf{x}_t, \mathbf{u}_t, \mathbf{w}_{t+1}) + K(\mathbf{x}_T) \right]$$

$$\sigma(\mathbf{u}_0) \subset \sigma(\mathbf{x}_0, \mathbf{w}_1), \dots, \sigma(\mathbf{u}_{T-1}) \subset \sigma(\mathbf{x}_0, \mathbf{w}_1, \dots, \mathbf{w}_T)$$

$$\mathbf{x}_{t+1} = f_t(\mathbf{x}_t, \mathbf{u}_t, \mathbf{w}_{t+1}), \quad t = 0, \dots, T-1$$

$$\mathbf{u}_t \in \mathbb{B}_t(\mathbf{x}_t, \mathbf{w}_{t+1}), \quad t = 0, \dots, T-1$$

Bellman equation and optimal policies in the hazard-decision information pattern

The uncertainty is observed before making the decision

```
initialization  $V_T = K$  ;  
for  $t = T, T - 1, \dots, 0$  do  
  forall  $x \in \mathbb{X}$  do  
    forall  $w \in \mathbb{W}_{t+1}$  do  
      forall  $u \in \mathbb{B}_t(x, w)$  do  
         $l_t(x, u, w) = L_t(x, u, w) + V_{t+1}(f_t(x, u, w))$   
         $\min_{u \in \mathbb{B}_t(x)} l_t(x, u, w)$  ;  
         $\mathbb{B}_t^*(x, w) = \arg \min_{u \in \mathbb{B}_t(x, w)} l_t(x, u, w)$   
       $V_t(x) = \sum_{w \in \mathbb{W}_{t+1}} \mathbb{P}\{w\} \min_{u \in \mathbb{B}_t(x, w)} l_t(x, u, w)$ 
```

When spilling decisions are made after knowing the water inflows, we obtain a linear dynamical model

$$\underbrace{S_{t+1}}_{\text{future volume}} = \underbrace{S_t}_{\text{volume}} - \underbrace{Q_t}_{\text{turbined}} + \underbrace{A_{t+1}}_{\text{inflow volume}} - \underbrace{R_{t+1}}_{\text{spilled}}$$

- ▶ S_t volume (stock) of water at the beginning of period $[t, t + 1[$
- ▶ A_{t+1} , inflow water volume (rain, etc.) during $[t, t + 1[$;
- ▶ Q_t turbined outflow volume
 - ▶ decided at the beginning of period $[t, t + 1[$ (hazard follows decision)
 - ▶ supposed to depend on the stock S_t
- ▶ R_{t+1} spilled volume
 - ▶ decided at the end of period $[t, t + 1[$ (hazard precedes decision)
 - ▶ supposed to depend on the stock S_t and on the inflow water A_t

$$0 \leq Q_t \leq S_t \text{ and } 0 \leq S_t - Q_t + A_{t+1} - R_{t+1} \leq S^\#$$

In the linear-convex case, value functions are convex

Here, we aim at **minimizing** expected cumulated costs

$$\mathbb{E} \left[\underbrace{\sum_{t=0}^{T-1} L_t(\mathbf{X}_t, \mathbf{U}_t, \mathbf{W}_{t+1})}_{\text{instantaneous cost}} + \underbrace{K(\mathbf{X}_T)}_{\text{final cost}} \right]$$

The value functions $x \mapsto V_t(x)$ are **convex** whenever

- ▶ the **instantaneous cost functions** $(x, u) \mapsto L_t(x, u, w)$ is **jointly convex in state and control** ($\forall w$)
- ▶ the **final cost function** $x \mapsto K(x)$ is **convex** ($\forall w$)
- ▶ the dynamics mappings are affine in state and control ($\forall w$)

$$f_t(x, u, w) = F_t(w)x + G_t(w)u + H_t(w)$$

- ▶ The constraint sets $\{(x, u) \mid u \in \mathbb{B}_t(x)\}$ are **convex**

The minimum over one variable of a jointly convex function is convex in the other variable

A lemma in convex analysis

Let $f : \mathbb{X} \times \mathbb{Y} \rightarrow \mathbb{R}$ be convex, and let $C \subset \mathbb{X} \times \mathbb{Y}$ be a convex set
Then, the so-called **marginal function** $g : \mathbb{X} \rightarrow \mathbb{R}$ defined by

$$g(x) = \min_{y \in \mathbb{Y}, (x,y) \in C} f(x,y), \quad \forall x \in \mathbb{X}$$

is a convex function

Stochastic Dual Dynamic Programming (SDDP)

- ▶ The dynamic programming equation associated with the problem of **minimizing the expected costs** is

$$V_T(x) = \overbrace{K(x)}^{\text{final cost}}$$
$$V_t(x) = \min_{u \in \mathbb{B}_t(x)} \mathbb{E}_{\mathbf{W}_{t+1}} \left[\overbrace{L_t(x, u, \mathbf{W}_{t+1})}^{\text{instantaneous cost}} + \underbrace{V_{t+1}(F_t(\mathbf{W}_{t+1})x + G_t(\mathbf{W}_{t+1})u + H_t(\mathbf{W}_{t+1}))}_{\text{future state}} \right]$$

- ▶ It can be shown by induction that $x \mapsto V_t(x)$ is convex
- ▶ A subgradient at x_{t+1}^* defines a hyperplane, hence a **lower affine approximation of the value function**, calculated by duality

SDDP and autoregressive noise models

The property that value functions are convex extends to the following cases

- ▶ Multiple stocks interconnected by linear dynamics

$$\mathbf{S}_{t+1}^i = \mathbf{S}_t^i + \mathbf{A}_t^i + \mathbf{Q}_t^{i-1} - \mathbf{Q}_t^i - \mathbf{R}_{t+1}^i$$

- ▶ Water inflows following an autoregressive model

$$\mathbf{A}_{t+1}^i = \sum_{k=0, \dots, K^i} \alpha_k \mathbf{A}_{t-k}^i + \mathbf{W}_{t+1}$$

where the random variables $\mathbf{W}_1, \dots, \mathbf{W}_{T-1}, \mathbf{W}_T$ are independent

Summary

- ▶ Bellman's Principle of Optimality breaks an intertemporal optimization problem into a sequence of **interconnected static optimization problems**
- ▶ The cost-to-go / value function / Bellman function is solution of a backward **dynamic programming equation**, or Bellman equation
- ▶ The Bellman equation provides an **optimal policy**, a concept of solution adapted to uncertain case
- ▶ In numerical practice, the **curse of dimensionality** forbids to use dynamic programming for a state with dimension more than four or five
- ▶ However, special cases like the linear-quadratic or the linear-convex ones, do not (totally) suffer from the curse of dimensionality