Single and Multi Agent Optimization, Game Theory with Information

Michel De Lara
Cermics, École des Ponts ParisTech
France

École des Ponts ParisTech

December 8, 2016

Outline of the presentation

One agent, one criterion optimization

Multi criteria optimization

Multiple agents: Witsenhausen intrinsic model

Team and dynamic optimization with information

Non-cooperative game theory with information

Bird's eye view from optimization to game theory

- Optimization

$$j : \mathbb{U} \to \mathbb{R}$$

- Multicriteria

$$j_a : \mathbb{U} \to \mathbb{R} \; , \quad a \in \mathbb{A}$$

- Non-cooperative game theory

$$j_a : \prod_{b \in \mathbb{A}} \mathbb{U}_b \to \mathbb{R} \; , \quad a \in \mathbb{A}$$

- Cooperative game theory

$$j : 2^{\mathbb{A}} \to \mathbb{R}$$

Outline of the presentation

One agent, one criterion optimization

Multi criteria optimization

Multiple agents: Witsenhausen intrinsic model

Team and dynamic optimization with information

Non-cooperative game theory with information

Let us start by lining up the ingredients for a general abstract optimization problem

- Optimization set $\mathbb{U}$ containing optimization variables $u \in \mathbb{U}$
- A criterion $J : \mathbb{U} \to \mathbb{R} \cup \{+\infty\}$
- Constraints of the form $u \in \mathbb{U}^{ad} \subset \mathbb{U}$

$$\min_{u \in \mathbb{U}^{ad}} J(u)$$

Outline of the presentation

One agent, one criterion optimization
- Deterministic optimization
- Optimization under uncertainty

Multi criteria optimization

Multiple agents: Witsenhausen intrinsic model

Team and dynamic optimization with information
- Team optimization
- Sequential (dynamic) stochastic optimization
- Special classical cases: SP and SOC

Non-cooperative game theory with information

Examples of classes of deterministic optimization problems

$$\min_{u \in \mathbb{U}^{ad}} J(u)$$

- Linear programming
  - Optimization set $\mathbb{U} = \mathbb{R}^N$
  - Criterion J is linear (affine)
  - Constraints $\mathbb{U}^{ad}$ defined by
    a finite number of linear (affine) equalities and inequalities
- Convex optimization
  - Criterion J is a convex function
  - Constraints $\mathbb{U}^{ad}$ define a convex set
- Combinatorial optimization
  - Optimization set $\mathbb{U}$ is discrete (binary $\{0, 1\}^N$, integer $\mathbb{Z}^N$, etc.)

A deterministic sequential optimization problem is just defined over a product space, without arrow of time

- A set $\{t_0, t_0+1, \ldots, T\} \subset \mathbb{N}$ of discrete times t
- Control sets $\mathbb{U}_t$ containing control variable $u_t \in \mathbb{U}_t$, for $t = t_0, t_0+1, \ldots, T$
- A criterion $J : \prod_{t=t_0}^{T} \mathbb{U}_t \to \mathbb{R} \cup \{+\infty\}$
- Constraints of the form $u = (u_{t_0}, \ldots, u_T) \in \mathbb{U}^{ad} \subset \prod_{t=t_0}^{T} \mathbb{U}_t$

$$\min_{(u_{t_0}, \ldots, u_T) \in \mathbb{U}^{ad}} J(u_{t_0}, \ldots, u_T)$$

Two-stage problem

Times $t \in \{0, 1\}$ (and criterion $L_0(u_0) + L_1(u_1, \omega)$)

Outline of the presentation

What makes optimization under uncertainty specific

- Optimization set is made of random variables
- Criterion generally derives from a mathematical expectation, or from a risk measure
- Constraints
  - generally include measurability constraints, like the nonanticipativity constraints,
  - and may also include probability constraints, or robust constraints

Here are the ingredients for a general abstract optimization problem under uncertainty

- A set $\mathbb{U}$
- A set Ω of scenarios
- An optimization set $\mathbb{V} \subset \mathbb{U}^{\Omega}$ containing random variables $\mathbf{V} : \Omega \to \mathbb{U}$
- A criterion $J : \mathbb{V} \to \mathbb{R} \cup \{+\infty\}$
- Constraints of the form $\mathbf{V} \in \mathbb{V}^{ad} \subset \mathbb{V}$

$$\min_{\mathbf{V} \in \mathbb{V}^{ad}} J(\mathbf{V})$$

Here is the most common framework for robust and stochastic optimization

- A set $\mathbb{U}$
- A set Ω of scenarios, or states of Nature, possibly equipped with a σ-algebra
- An optimization set $\mathbb{V} \subset \mathbb{U}^{\Omega}$ containing random variables $\mathbf{V} : \Omega \to \mathbb{U}$
- A risk measure $\mathbb{F} : \mathbb{V} \to \mathbb{R} \cup \{+\infty\}$
- A function $j : \mathbb{U} \times \Omega \to \mathbb{R} \cup \{+\infty\}$ (say, the "deterministic" criterion)
- Constraints of the form $\mathbf{V} \in \mathbb{V}^{ad} \subset \mathbb{V}$

$$\min_{\mathbf{V} \in \mathbb{V}^{ad}} J(\mathbf{V}) = \mathbb{F}\big[j(\mathbf{V}(\cdot), \cdot)\big]$$

where the notation means that the risk measure $\mathbb{F}$ has for argument the random variable

$$j(\mathbf{V}(\cdot), \cdot) : \Omega \to \mathbb{R} \cup \{+\infty\} \ , \ \ \omega \mapsto j(\mathbf{V}(\omega), \omega)$$

Examples of classes of robust and stochastic optimization problems

- Stochastic optimization "à la" gradient stochastique
  - The risk measure $\mathbb{F}$ is a mathematical expectation $\mathbb{E}$
  - Measurability constraints make that random variables $\mathbf{V} \in \mathbb{V}^{ad}$ are constant, that is, are deterministic decision variables

$$\min_{u \in \mathbb{U}^{ad}} \mathbb{E}_{\mathbb{P}}\big[j(u, \cdot)\big]$$

- Robust optimization
  - The risk measure $\mathbb{F}$ is the fear operator/worst case $\max_{\omega \in \overline{\Omega}}$, where $\overline{\Omega} \subset \Omega$
  - Measurability constraints make that random variables $\mathbf{V} \in \mathbb{V}^{ad}$ are constant, that is, are deterministic decision variables

$$\min_{u \in \mathbb{U}^{ad}} \max_{\omega \in \overline{\Omega}} j(u, \cdot)$$

Examples

- A set $\mathbb{U}$
  $\mathbb{U} = \mathbb{U}_0 \times \mathbb{U}_1$ in two stage programming
- A set Ω of scenarios
  Ω finite, $\Omega = \mathbb{N} \times \mathbb{W}^{\mathbb{N}}$ for discrete time stochastic processes
- An optimization set $\mathbb{V} \subset \mathbb{U}^{\Omega}$ containing random variables $\mathbf{V} : \Omega \to \mathbb{U}$
- A risk measure $\mathbb{F} : \mathbb{V} \to \mathbb{R} \cup \{+\infty\}$
  most often a mathematical expectation $\mathbb{E}$,
  but can be $\max_{\omega \in \overline{\Omega}}$ in the robust case, with $\overline{\Omega} \subset \Omega$
- A function $j : \mathbb{U} \times \Omega \to \mathbb{R} \cup \{+\infty\}$
- Constraints of the form $\mathbf{V} \in \mathbb{V}^{ad} \subset \mathbb{V}$
  - Measurability constraints,
    like the nonanticipativity constraints
  - Pointwise constraints,
    like probability constraints and robust constraints

Most common constraints in robust and stochastic optimization problems

- Measurability constraints

$$\mathbf{V} \in \text{linear subspace of } \mathbb{U}^{\Omega}$$

  like the nonanticipativity constraints $\mathbf{V} = (\mathbf{V}_0, \mathbf{V}_1)$,
  $\mathbf{V}_0$ is $\mathcal{F}_0$-measurable, $\mathbf{V}_1$ is $\mathcal{F}_1$-measurable
- Pointwise constraints, with $\mathbb{U}^{ad} : \Omega \rightrightarrows \mathbb{U}$
  - probability constraints

$$\mathbb{P}\big(\mathbf{V} \in \mathbb{U}^{ad}\big) \geq 1 - \epsilon$$

  - robust constraints

$$\mathbf{V}(\omega) \in \mathbb{U}^{ad}(\omega) \ , \ \ \forall \omega \in \overline{\Omega} \subset \Omega$$

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

The regret performs an additive normalization of the function $j : \mathbb{U} \times \Omega \to \mathbb{R} \cup \{+\infty\}$

Regret

For $u \in \mathbb{U}$ and $\omega \in \Omega$, the regret is

$$r(u, \omega) = j(u, \omega) - \min_{u' \in \mathbb{U}} j(u', \omega)$$

Then, take any risk measure $\mathbb{F}$ and solve

$$\min_{\mathbf{V} \in \mathbb{V}^{ad}} \mathbb{F}\big[r(\mathbf{V}, \cdot)\big] = \min_{\mathbf{V} \in \mathbb{V}^{ad}} \mathbb{F}\big[j(\mathbf{V}(\omega), \omega) - \min_{u \in \mathbb{U}} j(u, \omega)\big]$$

so that one can have minimal worse regret, minimal expected regret, etc.

Where have we gone till now? And what comes next

- A single criterion
- A single agent with all the information at hand
  (this is going to change in multi-agent optimization)

Outline of the presentation

One agent, one criterion optimization

Multi criteria optimization

Multiple agents: Witsenhausen intrinsic model

Team and dynamic optimization with information

Non-cooperative game theory with information

Here are the ingredients for a multi criteria optimization problem

- A set $\mathbb{U}$
- A finite set $\mathbb{A}$ of stake holders
- A collection of criteria $J_a : \mathbb{U} \to \mathbb{R} \cup \{+\infty\}$, for $a \in \mathbb{A}$

In multi criteria optimization, stake holders $a \in \mathbb{A}$ bargain over a common decision $u \in \mathbb{U}$

In a multi criteria optimization problem, a solution is a Pareto optimum

A decision $u^\flat \in \mathbb{U}$ is dominated by a decision $u^\sharp \in \mathbb{U}$ if

- all stake holders prefer $u^\sharp$ to $u^\flat$, that is,

$$J_a(u^\sharp) \geq J_a(u^\flat) \ , \quad \forall a \in \mathbb{A}$$

- at least one stake holder strictly prefers $u^\sharp$ to $u^\flat$, that is,

$$\exists a \in \mathbb{A} \ , \quad J_a(u^\sharp) > J_a(u^\flat)$$

A decision is a Pareto optimum if it is not dominated by any other decision

Outline of the presentation

One agent, one criterion optimization

Multi criteria optimization

Multiple agents: Witsenhausen intrinsic model

Team and dynamic optimization with information

Non-cooperative game theory with information

Witsenhausen intrinsic model

Till now, we could only account for agents whose order of play was fixed in advance (sequential optimization)

To account for agents whose order of play is not fixed in advance,
but depends on the state of Nature and on the moves of other agents,
we use the Witsenhausen intrinsic model
with an information field attached to each agent

Outline of the presentation

One agent, one criterion optimization

Multi criteria optimization

Multiple agents: Witsenhausen intrinsic model

Team and dynamic optimization with information

Non-cooperative game theory with information

Outline of the presentation

Outline of the presentation

Let us line up the ingredients for a stochastic sequential optimization problem

- A set $\{t_0, t_0+1, \ldots, T\} \subset \mathbb{N}$ of discrete times, with generic element t
- Control sets $\mathbb{U}_t$ containing control variable $u_t \in \mathbb{U}_t$, for $t = t_0, t_0+1, \ldots, T$
- Constraints of the form $u_t \in \mathbb{U}_t^{ad} \subset \mathbb{U}_t$
- A set Ω of scenarios, or states of Nature, with generic element ω (without temporal structure, a priori)
- A pre-criterion $j : \mathbb{U}_{t_0} \times \cdots \times \mathbb{U}_T \times \Omega \to \mathbb{R}$, with generic value $j(u_{t_0}, \ldots, u_T, \omega)$

Two-stage problem

Times $t \in \{0, 1\}$ (and pre-criterion $L_0(u_0) + L_1(u_1, \omega)$)

- Stochastic optimization deals with risk attitudes:
  mathematical expectation $\mathbb{E}$, risk measure $\mathbb{F}$ (including worst case),
  probability or robust constraints
- Stochastic dynamic optimization emphasizes
  the handling of online information,
  and especially the nonanticipativity constraints

For the purpose of handling online information, we introduce fields and subfields

1. $(\Omega, \mathcal{F})$ a measurable space (uncertainties, states of Nature)
2. $(\mathbb{U}_{t_0}, \mathcal{U}_{t_0}), \ldots, (\mathbb{U}_T, \mathcal{U}_T)$ measurable spaces (decision spaces)
3. Subfield $\mathcal{I}_t \subset \mathcal{U}_{t_0} \otimes \cdots \otimes \mathcal{U}_{t-1} \otimes \mathcal{F}$, for $t = t_0, \ldots, T$ (information)

The inclusion

$$\underbrace{\mathcal{I}_t}_{\text{information}} \subset \underbrace{\mathcal{U}_{t_0} \otimes \cdots \otimes \mathcal{U}_{t-1}}_{\text{past controls}} \otimes \mathcal{F}$$

captures the fact that the information at time t is made at most of past controls and of the state of Nature (causality)

Static team

Subfield $\mathcal{I}_t \subset \mathcal{F}$ for $t = t_0, \ldots, T$ (no dynamic flow of information)

We introduce strategies

Decision rule, policy, strategy

A strategy is a sequence $\lambda = \{\lambda_t\}_{t=t_0,\ldots,T}$
of measurable mappings from past histories to decision sets

$$
\begin{aligned}
\lambda_{t_0} &: (\Omega, \mathcal{F}) \to (\mathbb{U}_{t_0}, \mathcal{U}_{t_0}) \\
&\ldots \\
\lambda_t &: (\mathbb{U}_{t_0} \times \cdots \times \mathbb{U}_{t-1} \times \Omega, \mathcal{U}_{t_0} \otimes \cdots \otimes \mathcal{U}_{t-1} \otimes \mathcal{F}) \to (\mathbb{U}_t, \mathcal{U}_t) \\
&\ldots
\end{aligned}
$$

With obvious notations, the set of strategies is denoted by

$$\Lambda_{t_0,\ldots,T} = \prod_{t=t_0,\ldots,T} \Lambda_t$$

We introduce admissible strategies to account for the interplay between decision and information

Admissible strategy

An admissible strategy is a strategy $\lambda = \{\lambda_t\}_{t=t_0,\dots,T}$

$$
\begin{aligned}
&\lambda_{t_0} : (\Omega, \mathcal{F}) \to (\mathbb{U}_{t_0}, \mathcal{U}_{t_0}) \\
&\dots \\
&\lambda_t : (\mathbb{U}_{t_0} \times \cdots \times \mathbb{U}_{t-1} \times \Omega, \mathcal{U}_{t_0} \otimes \cdots \otimes \mathcal{U}_{t-1} \otimes \mathcal{F}) \to (\mathbb{U}_t, \mathcal{U}_t) \\
&\dots
\end{aligned}
$$

satisfying, for $t = t_0, \dots, T$, the information constraints

$$
\lambda_t^{-1}(\mathcal{U}_t) \subset \underbrace{\mathcal{I}_t}_{\text{information}}
$$

With obvious notations, the set of admissible strategies is denoted by

$$
\Lambda^{ad}_{t_0,\dots,T} = \prod_{t=t_0,\dots,T} \Lambda^{ad}_t
$$

The solution map is attached to a strategy, and maps a scenario towards a history

Solution map

With a strategy λ, we associate the mapping

$$S_\lambda : \Omega \to \underbrace{\mathbb{U}_{t_0} \times \cdots \times \mathbb{U}_T \times \Omega}_{\text{history space}}$$

called solution map, and defined by

$$(u_{t_0}, \ldots, u_T, \omega) = S_\lambda(\omega) \iff \begin{cases} u_{t_0} & = \lambda_{t_0}(\omega) \\ u_{t_0+1} & = \lambda_{t_0+1}(u_{t_0}, \omega) \\ \vdots & \vdots \\ u_T & = \lambda_T(u_{t_0}, \cdots, u_{T-1}, \omega) \end{cases}$$

By composing the pre-criterion with the solution map, we move forward the design of a criterion

- With a strategy λ, we associate the solution map

$$S_\lambda : \Omega \to \underbrace{\mathbb{U}_{t_0} \times \cdots \times \mathbb{U}_T \times \Omega}_{\text{history space}}$$

  that maps a scenario towards a history
- The pre-criterion

$$j : \mathbb{U}_{t_0} \times \cdots \times \mathbb{U}_T \times \Omega \to \mathbb{R}$$

  maps a a history towards the real numbers
- Therefore, by composing the pre-criterion with the solution map, we obtain

$$j \circ S_\lambda : \Omega \to \mathbb{R}$$

  that maps a scenario towards the real numbers

For the purpose of building a criterion (and of handling risk attitudes), we introduce a risk measure

As $j \circ S_\lambda \in \mathbb{R}^\Omega$, all we need is a risk measure

$$\mathbb{F} : \mathbb{R}^\Omega \to \mathbb{R} \cup \{+\infty\}$$

to build a criterion that maps a strategy λ
towards the (extended) real numbers

$$\lambda \in \Lambda_{t_0,\ldots,T} \mapsto \mathbb{F} \circ j \circ S_\lambda \in \mathbb{R} \cup \{+\infty\}$$

where we recall that $\Lambda_{t_0,\ldots,T}$ denotes the set of strategies

We can now formulate an optimization problem under uncertainty

Optimization problem under uncertainty

When $\mathbb{F}$ is a risk measure on Ω,

$$\mathbb{F} : \mathbb{R}^{\Omega} \to \mathbb{R} \cup \{+\infty\} ,$$

the corresponding optimization problem under uncertainty is

$$\min_{\lambda \in \Lambda^{ad}_{t_0,\ldots,T}} \mathbb{F}\Big(j\big(S_{\lambda}(\cdot)\big) \Big)$$

where we recall that $\Lambda^{ad}_{t_0,\ldots,T}$ denotes the set of admissible strategies, those such that

$$\lambda_t^{-1}(\mathcal{U}_t) \subset \mathcal{I}_t \ , \ \ \forall t = t_0, \ldots, T$$

Risk neutral and robust optimization appear as special cases

Risk-neutral stochastic optimization problem

When $(\Omega, \mathcal{F}, \mathbb{P})$ is a probability space,
the stochastic optimization problem is

$$\min_{\lambda \in \Lambda^{ad}_{t_0,\dots,T}} \mathbb{E}_{\mathbb{P}}\Big(j\big(S_\lambda(\cdot)\big)\Big)$$

Robust optimization problem

When $\overline{\Omega} \subset \Omega$, the robust optimization problem is

$$\min_{\lambda \in \Lambda^{ad}_{t_0,\dots,T}} \max_{\omega \in \overline{\Omega}} j\big(S_\lambda(\omega)\big)$$

Outline of the presentation

One agent, one criterion optimization

- Deterministic optimization
- Optimization under uncertainty

Multi criteria optimization

Multiple agents: Witsenhausen intrinsic model

Team and dynamic optimization with information

- Team optimization
- Sequential (dynamic) stochastic optimization
- Special classical cases: SP and SOC

Non-cooperative game theory with information

- We survey special cases of nonanticipativity constraints, when the scenario space is a product over time
- We show how two classical settings fit in the geneal framework: stochastic programming (SP) and stochastic optimal control (SOC)

How to handle the nonanticipativity constraints

- Product scenario space

$$\Omega = \prod_{t=t_0+1}^{T} \mathbb{W}_t \text{ with } \mathcal{F} = \bigotimes_{t=t_0+1}^{T} \mathcal{W}_t$$

- Past uncertainties fields for $t = t_0 + 1, \ldots, T$,

$$\mathcal{F}_t = \underbrace{\mathcal{W}_{t_0+1} \otimes \cdots \otimes \mathcal{W}_t}_{\text{past uncertainties}} \otimes \{\emptyset, \mathbb{W}_{t+1}\} \otimes \cdots \otimes \{\emptyset, \mathbb{W}_T\}$$

- Nonanticipativity constraint

$$\mathcal{I}_{t_0} = \{\emptyset, \Omega\} \text{ and } \mathcal{I}_t \subset \underbrace{\mathcal{U}_{t_0} \otimes \cdots \otimes \mathcal{U}_{t-1}}_{\text{past controls}} \otimes \mathcal{F}_t$$

Two-stage stochastic programming problem

$$\min_{u_0} L_0(u_0) + \mathbb{E}\Big(\min_{u_1} L_1(u_1, \omega_1) \Big)$$

- Decision spaces

$$(\mathbb{U}_0, \mathcal{U}_0) = (\mathbb{R}^{p_0}, \mathcal{B}^{o}_{\mathbb{R}^{p_0}}) \text{ and } (\mathbb{U}_1, \mathcal{U}_1) = (\mathbb{R}^{p_1}, \mathcal{B}^{o}_{\mathbb{R}^{p_1}})$$

- Probability $\mathbb{P}$ on the probability space

$$\Omega = \mathbb{W}_1 = \mathbb{R}^{q_1} \text{ with } \mathcal{F} = \mathcal{B}^{o}_{\mathbb{W}_1} = \mathcal{B}^{o}_{\mathbb{R}^{q_1}}$$

- Information fields

$$\mathcal{I}_0 = \{\emptyset, \Omega\} \text{ and } \mathcal{I}_1 = \mathcal{U}_0 \otimes \mathcal{F}$$

  - at the first stage, there is no information whatsoever
  - at the second stage, the first decision and the state of Nature are available for decision-making

Multi-stage stochastic programming problem

$$\min_{u_{t_0}} L_{t_0}(u_{t_0}) + \mathbb{E}\Big(\min_{u_{t_0+1}} L_{t_0+1}(u_{t_0+1}, \omega_{t_0+1}) + \mathbb{E}\Big(\cdots + \mathbb{E}\Big(\min_{u_T} L_T(u_T, \omega_T) \Big)\Big)\Big) ,$$

This corresponds to the decision spaces

$$(\mathbb{U}_{t_0}, \mathcal{U}_{t_0}) = (\mathbb{R}^{p_{t_0}}, \mathcal{B}^{o}_{\mathbb{R}^{p_{t_0}}}) , \quad \ldots , \quad (\mathbb{U}_T, \mathcal{U}_T) = (\mathbb{R}^{p_T}, \mathcal{B}^{o}_{\mathbb{R}^{p_T}}) ,$$

and to the probability space

$$\Omega = \prod_{t=t_0+1}^{T} \mathbb{W}_t \text{ with } \mathcal{F} = \bigotimes_{t=t_0+1}^{T} \mathcal{W}_t$$

equipped with a probability $\mathbb{P}$

State model and stochastic optimal control (SOC)

- Dynamics with an intermediary variable $x_t \in \mathbb{X}_t$

$$x_{t+1} = f_t\big(x_t, u_t, w_t\big) \ , \quad t = t_0, \ldots, T$$

- Criterion $j\big(x(\cdot), u(\cdot), w(\cdot)\big)$ defined over trajectories
- White noise assumption: the scenario space $\Omega = \prod_{t=t_0+1}^{T} \mathbb{W}_t$ is equipped with a product probability

$$\mathbb{P} = \bigotimes_{t=t_0+1}^{T} \mu_t$$

- Then $x_t \in \mathbb{X}_t$ deserves the name of state:
  x_t summarizes the past history in that, given time t and the value of x_t, one can calculate the optimal u_t and also x_{t+1} without knowledge of the whole history $(u_{t_0}, \ldots, u_{t-1}, \omega)$, for all t

Where have we gone till now? And what comes next

- A single criterion
  (this is going to change in game theory)
- Multiple agents with different information

Outline of the presentation

One agent, one criterion optimization

Multi criteria optimization

Multiple agents: Witsenhausen intrinsic model

Team and dynamic optimization with information

Non-cooperative game theory with information