Single and Multi Agent Optimization, Game Theory with Information

Michel De Lara
Cermics, École des Pontes ParisTech
France

École des Ponts ParisTech

November 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

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 $L_0\big(u_0\big) + L_1\big(u_1, \omega\big)$)

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

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

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 $L_0\big(u_0\big) + L_1\big(u_1, \omega\big)\big)$

- 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. $(\mathbb{U}_{t_0}, \mathcal{U}_{t_0})$,..., $(\mathbb{U}_T, \mathcal{U}_T)$ 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 $\lambda = \{\lambda_t\}_{t=t_0,\dots,T}$
of measurable mappings from past histories to decision sets

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

With obvious notations, the set of strategies is denoted by

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

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

$$\Lambda_{t_0,\dots,T}^{ad} = \prod_{t=t_0,\dots,T} \Lambda_t^{ad}$$

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

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

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

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

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

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

Robust optimization problem

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

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

Outline of the presentation

- 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 $t = t_0 + 1, \dots, T$,

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

- 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

Outline of the presentation

Witsenhausen intrinsic model with Nature and two players

We lay out

- basic sets
  - decision sets
  - states of Nature
  - history set

  and their σ-fields
- objective functions
- beliefs
- information σ-fields, admissible strategies and predecessors

Nature moves and players decisions

- Let Ω be a measurable set equipped with σ-field $\mathcal{F}$ which represents all uncertainties:
  any $\omega \in \Omega$ is called a state of Nature
- The player a makes one decision $u_a \in \mathbb{U}_a$
  where the decision set $\mathbb{U}_a$ is equipped with a σ-field $\mathcal{U}_a$
- The player b makes one decision $u_b \in \mathbb{U}_b$
  where the decision set $\mathbb{U}_b$ is equipped with a σ-field $\mathcal{U}_b$

History space

The history space is the product space

$$\mathbb{H} = \mathbb{U}_a \times \mathbb{U}_b \times \Omega$$

equipped with the product history field

$$\mathcal{H} = \mathcal{U}_a \otimes \mathcal{U}_b \otimes \mathcal{F}$$

Criteria (costs or payoffs)

- For any history $h = (u_a, u_b, \omega) \in \mathbb{H}$,
  player a and player b undergo costs or payoffs
- Costs or payoffs are materialized under the form of (measurable) objective functions or criteria

$$j_a : \mathbb{U}_a \times \mathbb{U}_b \times \Omega \to \mathbb{R}$$

$$j_b : \mathbb{U}_a \times \mathbb{U}_b \times \Omega \to \mathbb{R}$$

Beliefs (and risk attitudes)

- All sources of randomness are collected in the set Ω
- Player a and player b express beliefs and risk attitudes towards events in $\mathcal{F}$
- We denote real-valued random variables on $(\Omega, \mathcal{F})$ by

$$\mathbb{L}(\Omega, \mathcal{F}) = \{\mathbf{X} : (\Omega, \mathcal{F}) \to (\mathbb{R}, \mathcal{B}_{\mathbb{R}}) \ , \ \ \mathbf{X}^{-1}(\mathcal{B}_{\mathbb{R}}) \subset \mathcal{F}\}$$

Belief

A belief of the player $c \in \{a, b\}$ is
a probability distribution $\mathbb{P}_c$ over $(\Omega, \mathcal{F})$

- We denote by

$$\mathbb{L}(\Omega, \mathcal{F}, \mathbb{P}_c) = \{\mathbf{X} \in \mathbb{L}(\Omega, \mathcal{F}) \ , \ \ \mathbb{E}_{\mathbb{P}_c}\big(|\mathbf{X}|\big) < +\infty\}$$

the space of $\mathbb{P}_c$-integrable random variables,
where $\mathbb{E}_{\mathbb{P}_c}$ denotes the mathematical expectation

Information, admissible strategies and predecessors

Information

- When making a decision,
  player a and player b can make use of information,
  materialized under the form of σ-fields
- The information fields $\mathcal{I}_a$ of the player a and $\mathcal{I}_b$ of the player b are subfields of $\mathcal{H}$

$$\mathcal{I}_a \subset \mathcal{U}_a \otimes \mathcal{U}_b \otimes \mathcal{F}$$
$$\mathcal{I}_b \subset \mathcal{U}_a \otimes \mathcal{U}_b \otimes \mathcal{F}$$

Classical information patterns in game theory

Two players: the principal P (leader) and the agent A (follower)

- Stackelberg leadership model

$$\mathcal{I}_A \subset \{\emptyset, \mathbb{U}_A\} \otimes \mathcal{U}_P \otimes \mathcal{F} \,, \;\; \mathcal{I}_P \subset \{\emptyset, \mathbb{U}_A\} \otimes \{\emptyset, \mathbb{U}_P\} \otimes \mathcal{F}$$

- Moral hazard

$$\mathcal{I}_P \subset \{\emptyset, \mathbb{U}_A\} \otimes \{\emptyset, \mathbb{U}_P\} \otimes \mathcal{F}$$

- Adverse selection

$$\{\emptyset, \mathbb{U}_A\} \otimes \{\emptyset, \mathbb{U}_P\} \otimes \mathcal{F} \subset \mathcal{I}_A \,, \;\; \mathcal{I}_P \subset \mathcal{U}_A \otimes \{\emptyset, \mathbb{U}_P\} \otimes \{\emptyset, \Omega\}$$

- Signaling

$$\{\emptyset, \mathbb{U}_A\} \otimes \{\emptyset, \mathbb{U}_P\} \otimes \mathcal{F} \subset \mathcal{I}_A \,, \;\; \mathcal{I}_P = \mathcal{U}_A \otimes \{\emptyset, \mathbb{U}_P\} \otimes \{\emptyset, \Omega\}$$

Admissible strategies

- Information is the fuel of strategies
- A strategy of the player $c \in \{a, b\}$ is a mapping

$$\lambda_c : \mathbb{U}_a \times \mathbb{U}_b \times \Omega \to \mathbb{U}_c$$

Admissible strategy

An admissible strategy of the player $c \in \{a, b\}$ is a mapping

$$\lambda_c : \mathbb{U}_a \times \mathbb{U}_b \times \Omega \to \mathbb{U}_c \text{ such that } \sigma^{-1}(\mathcal{U}_c) \subset \mathcal{I}_c$$

- The set of admissible strategies of the player c is

$$\Lambda_c^{ad} = \{\lambda_c \mid \mathbb{U}_a \times \mathbb{U}_b \times \Omega \to \mathbb{U}_c \ , \ \ \sigma^{-1}(\mathcal{U}_c) \subset \mathcal{I}_c\}$$

- The set of admissible strategies is

$$\Lambda^{ad} = \Lambda_a^{ad} \times \Lambda_b^{ad}$$

Absence of "self-information" and structure of admissible strategies

- The information fields $\mathcal{I}_a$ and $\mathcal{I}_b$ display the absence of "self-information" when

$$\mathcal{I}_a \subset \mathcal{U}_b \otimes \{\emptyset, \mathbb{U}_a\} \otimes \mathcal{F}$$

$$\mathcal{I}_b \subset \{\emptyset, \mathbb{U}_b\} \otimes \mathcal{U}_a \otimes \mathcal{F}$$

- When σ-fields include singletons and we exclude "self-information", then, for any admissible strategy λ_c of the player $c \in \{a, b\}$, we have that the expression $\lambda_c(u_a, u_b, \omega)$ does not depend on u_c:

$$\lambda_a(\cancel{u_a}, u_b, \omega) = \widetilde{\lambda_a}(u_b, \omega) \; , \;\; \lambda_b(u_a, \cancel{u_b}, \omega) = \widetilde{\lambda_b}(u_a, \omega)$$

Solvability property

The information fields $\mathcal{I}_a$ and $\mathcal{I}_b$ display the solvability property when,

- for any couple $(\lambda_a, \lambda_b) \in \Lambda_a^{ad} \times \Lambda_b^{ad}$ of admissible strategies and any state of Nature $\omega \in \Omega$,
- there exists one, and only one, couple $(u_a, u_b) \in \mathbb{U}_a \times \mathbb{U}_b$ of decisions such that

$$u_a = \lambda_a(u_a, u_b, \omega)$$

$$u_b = \lambda_b(u_a, u_b, \omega)$$

Solvability property and solution map

Solution map

In case of solvability, we can define $S_{(\lambda_a,\lambda_b)}(\omega)$, for any $\omega \in \Omega$, by

$$S_{(\lambda_a,\lambda_b)}(\omega) = (u_a, u_b, \omega) \iff \begin{cases} u_a = \lambda_a(u_a, u_b, \omega) \\ u_b = \lambda_b(u_a, u_b, \omega) \end{cases}$$

Hence, we obtain a mapping called the solution map

$$S_{(\lambda_a,\lambda_b)} : \Omega \to \mathbb{H}$$

- The solvability property holds true in the sequential cases
- The graph of $S_{(\lambda_a,\lambda_b)}$ belongs to $\mathcal{I}_a \vee \mathcal{U}_a \vee \mathcal{I}_b \vee \mathcal{U}_b$.

Definition of predecessor, excluding Nature

Consider a subset B of $\{a, b\}$ — $B \in \{\emptyset, \{a\}, \{b\}, \{a, b\}\}$ — and define

$$\mathcal{H}_B = \prod_{c \in B} \mathcal{U}_c \otimes \prod_{c \notin B} \{\emptyset, \mathbb{U}_c\} \otimes \mathcal{F}$$

Predecessor

For any player $c \in \{a, b\}$, we define $\langle c \rangle_{\mathfrak{P}}$
as the intersection of all subsets B of $\{a, b\}$ such that $\mathcal{I}_c \subset \mathcal{H}_B$

$$\langle c \rangle_{\mathfrak{P}} = \bigcap_{B, \ \mathcal{I}_c \subset \mathcal{H}_B} B$$

When non empty, an element of $\langle c \rangle_{\mathfrak{P}}$ is called a predecessor of c

- Nature has no predecessor: Nature plays before the players (but is not necessarily revealed to the players)
- As an illustration, absence of "self-information" is equivalent to $c \notin \langle c \rangle_{\mathfrak{P}}$, for any $c \in \{a, b\}$

Common information patterns

- Sequential patterns
  - When $\langle a \rangle_{\mathfrak{P}} = \emptyset$ and $\langle b \rangle_{\mathfrak{P}} = \{a\}$,
    player a plays first, player b plays second
  - When $\langle b \rangle_{\mathfrak{P}} = \emptyset$ and $\langle a \rangle_{\mathfrak{P}} = \{b\}$,
    player b plays first, player a plays second
  - When $\langle a \rangle_{\mathfrak{P}} = \emptyset$ and $\langle b \rangle_{\mathfrak{P}} = \emptyset$,
    player a and player b play first (static team)
- Deadlock pattern
  - When $\langle a \rangle_{\mathfrak{P}} = \{b\}$ and $\langle b \rangle_{\mathfrak{P}} = \{a\}$,
    player a and player b are in a deadlock (non causal system)

Co-cycle property of the solution map (I)

- We suppose that $\langle b \rangle_{\mathfrak{P}} = \emptyset$ and $\langle a \rangle_{\mathfrak{P}} = \{b\}$,
  that is, player b plays first, player a plays second
- We consider a couple $(\lambda_a, \lambda_b) \in \Lambda_a^{ad} \times \Lambda_b^{ad}$ of admissible strategies

Co-cycle property of the solution map

We have that

- the strategy λ_b can be identified with $\lambda_b : \Omega \to \mathbb{U}_b$ and the partial solution map $S_{\lambda_b} : \Omega \to \mathbb{U}_b \times \Omega$ is such that $S_{\lambda_b}(\omega) = (\lambda_b(\omega), \omega)$
- the strategy λ_a can be identified with $\lambda_a : \mathbb{U}_b \times \Omega \to \mathbb{U}_a$
- the solution map has the following co-cycle property

$$S_{(\lambda_a,\lambda_b)} = (\lambda_a \circ S_{\lambda_b}, S_{\lambda_b}) : \Omega \to \mathbb{U}_a \times \big(\mathbb{U}_b \times \Omega\big)$$

$$S_{(\lambda_a,\lambda_b)}(\omega) = \Big(\lambda_a\Big(\lambda_b(\omega), \omega\Big), \lambda_b(\omega), \omega\Big) \ , \quad \forall \omega \in \Omega$$

Co-cycle property of the solution map (II)

The co-cycle property

$$S_{(\lambda_a,\lambda_b)} = (\lambda_a \circ S_{\lambda_b}, S_{\lambda_b})$$

is equivalent to

$$S_{(\lambda_a,\lambda_b)}(\omega) = (u_a, u_b, \omega) \iff \begin{cases} (u_b, \omega) & = S_{\lambda_b}(\omega) \\ u_a & = \lambda_a(u_b, \omega) \end{cases}$$

Outline of the presentation

Criteria composed with solution map

- Costs or payoffs are

$$j_a : \mathbb{H} \to \mathbb{R}$$

$$j_b : \mathbb{H} \to \mathbb{R}$$

- Solution map is

$$S_{(\lambda_a,\lambda_b)} : \Omega \to \mathbb{H}$$

- The composition of criteria with the solution map provides random variables

$$j_a \circ S_{(\lambda_a,\lambda_b)} : \Omega \to \mathbb{R}$$

$$j_b \circ S_{(\lambda_a,\lambda_b)} : \Omega \to \mathbb{R}$$

Nash equilibrium for all states of Nature

Nash equilibrium for all states of Nature

We say that the couple $(\overline{\lambda}_a, \overline{\lambda}_b) \in \Lambda_a^{ad} \times \Lambda_b^{ad}$ of admissible strategies is a Nash equilibrium for all states of Nature if (in case of payoffs)

$$j_a \circ S_{(\overline{\lambda}_a, \overline{\lambda}_b)}(\omega) \geq j_a \circ S_{(\lambda_a, \overline{\lambda}_b)}(\omega) \ , \ \ \forall \lambda_a \in \Lambda_a^{ad} \ , \ \ \forall \omega \in \Omega$$

$$j_b \circ S_{(\overline{\lambda}_a, \overline{\lambda}_b)}(\omega) \geq j_b \circ S_{(\overline{\lambda}_a, \lambda_b)}(\omega) \ , \ \ \forall \lambda_b \in \Lambda_b^{ad} \ , \ \ \forall \omega \in \Omega$$

Bayesian Nash equilibrium

We suppose that player a has belief $\mathbb{P}_a$ and player b has belief $\mathbb{P}_b$

Bayesian Nash equilibrium

We say that the couple $(\overline{\lambda}_a, \overline{\lambda}_b) \in \Lambda_a^{ad} \times \Lambda_b^{ad}$ of admissible strategies is a Bayesian Nash equilibrium if (in case of payoffs)

$$\mathbb{E}_{\mathbb{P}_a}\left[j_a \circ S_{(\overline{\lambda}_a, \overline{\lambda}_b)}\right] \geq \mathbb{E}_{\mathbb{P}_a}\left[j_a \circ S_{(\lambda_a, \overline{\lambda}_b)}\right] , \quad \forall \lambda_a \in \Lambda_a^{ad}$$

$$\mathbb{E}_{\mathbb{P}_b}\left[j_b \circ S_{(\overline{\lambda}_a, \overline{\lambda}_b)}\right] \geq \mathbb{E}_{\mathbb{P}_b}\left[j_b \circ S_{(\overline{\lambda}_a, \lambda_b)}\right] , \quad \forall \lambda_b \in \Lambda_b^{ad}$$

Bayesian Nash equilibrium (with conditioning)

We suppose that player a has belief $\mathbb{P}_a$ and player b has belief $\mathbb{P}_b$

Bayesian Nash equilibrium (with conditioning)

We say that the couple $(\overline{\lambda}_a, \overline{\lambda}_b) \in \Lambda_a^{ad} \times \Lambda_b^{ad}$ of admissible strategies is a Bayesian Nash equilibrium (with conditioning) if (in case of payoffs)

$$\mathbb{E}_{\mathbb{P}_a}\left[j_a \circ S_{(\overline{\lambda}_a,\overline{\lambda}_b)} \mid S^{-1}_{(\overline{\lambda}_a,\overline{\lambda}_b)}(\mathcal{J}_a)\right] \geq \mathbb{E}_{\mathbb{P}_a}\left[j_a \circ S_{(\lambda_a,\overline{\lambda}_b)} \mid S^{-1}_{(\lambda_a,\overline{\lambda}_b)}(\mathcal{J}_a)\right] \quad \forall \lambda_a \in \Lambda_a^{ad}$$

$$\mathbb{E}_{\mathbb{P}_b}\left[j_b \circ S_{(\overline{\lambda}_a,\overline{\lambda}_b)} \mid S^{-1}_{(\overline{\lambda}_a,\overline{\lambda}_b)}(\mathcal{J}_b)\right] \geq \mathbb{E}_{\mathbb{P}_b}\left[j_b \circ S_{(\overline{\lambda}_a,\lambda_b)} \mid S^{-1}_{(\overline{\lambda}_a,\lambda_b)}(\mathcal{J}_b)\right] \quad \forall \lambda_b \in \Lambda_b^{ad}$$

Strategies have a dual impact on Bayesian Nash equilibrium (with conditioning)

- Consider agent a (the same analysis would hold for agent b)
- Observe that when agent a varies his strategy λ_a, the impact is dual:
  - the criterion is affected through $S_{(\overline{\lambda}_a,\overline{\lambda}_b)}$ in $j_a \circ S_{(\overline{\lambda}_a,\overline{\lambda}_b)}$
  - the information is affected through $S^{-1}_{(\lambda_a,\overline{\lambda}_b)}(\mathcal{J}_a)$

Bayesian Nash equilibrium
Equivalent formulation

- In fact, all changes in strategies (λ_a, λ_b) can be captured in the probability images

$$\mathbb{P}_a \circ S^{-1}_{(\lambda_a,\lambda_b)} \ , \ \ \mathbb{P}_b \circ S^{-1}_{(\lambda_a,\lambda_b)}$$

  on $(\mathbb{H}, \mathcal{H})$ of the probabilities $\mathbb{P}_a$ and $\mathbb{P}_b$ by the mappings $S_{(\lambda_a,\lambda_b)} : (\Omega, \mathcal{F}) \to (\mathbb{H}, \mathcal{H})$

- The couple $(\overline{\lambda}_a, \overline{\lambda}_b) \in \Lambda_a^{ad} \times \Lambda_b^{ad}$ of admissible strategies is a Bayesian Nash equilibrium if (in case of payoffs)

$$\mathbb{E}_{\mathbb{P}_a \circ S^{-1}_{(\overline{\lambda}_a,\overline{\lambda}_b)}} \Big[j_a \mid \mathcal{I}_a \Big] \geq \mathbb{E}_{\mathbb{P}_a \circ S^{-1}_{(\lambda_a,\overline{\lambda}_b)}} \Big[j_a \mid \mathcal{I}_a \Big] \ , \ \ \forall \lambda_a \in \Lambda_a^{ad}$$

$$\mathbb{E}_{\mathbb{P}_b \circ S^{-1}_{(\overline{\lambda}_a,\overline{\lambda}_b)}} \Big[j_b \mid \mathcal{I}_b \Big] \geq \mathbb{E}_{\mathbb{P}_b \circ S^{-1}_{(\overline{\lambda}_a,\lambda_b)}} \Big[j_b \mid \mathcal{I}_b \Big] \ , \ \ \forall \lambda_b \in \Lambda_b^{ad}$$

Outline of the presentation

Principal-agent models with two players

- A branch of Economics studies so-called principal-agent models
- Principal-agent models display a general information structure, which can be transparently expressed thanks to Witsenhausen intrinsic model
- The model exhibits two players
  - the principal P (leader), makes decisions $u_P \in \mathbb{U}_P$, where the set of decisions is equipped with a σ-field $\mathcal{U}_P$
  - the agent A (follower) makes decisions $u_A \in \mathbb{U}_A$, where the set of decisions is equipped with a σ-field $\mathcal{U}_A$
- and Nature, corresponding to private information (or type) of the agent A
  - Nature selects $\omega \in \Omega$, where Ω is equipped with a σ-field $\mathcal{F}$

Here is the most general information structure of principal-agent models

$$\mathfrak{I}_P \subset \mathfrak{U}_A \otimes \{\emptyset, \mathbb{U}_P\} \otimes \mathfrak{F}$$

$$\mathfrak{I}_A \subset \{\emptyset, \mathbb{U}_A\} \otimes \mathfrak{U}_P \otimes \mathfrak{F}$$

- By these expressions of the information fields
  - $\mathfrak{I}_P$ of the principal P (leader)
  - $\mathfrak{I}_A$ of the agent A (follower)
- we have excluded self-information, that is, we suppose that the information of a player cannot be influenced by his actions

Classical information patterns in game theory

Now, we will make the information structure more specific

- Stackelberg leadership model
- Moral hazard
- Adverse selection
- Signaling

Stackelberg leadership model

- In the Stackelberg leadership model of game theory,
- the follower A may partly observe the action of the leader P

$$\mathcal{I}_A \subset \{\emptyset, \mathbb{U}_A\} \otimes \mathcal{U}_P \otimes \mathcal{F}$$

- whereas the leader P observes at most the state of Nature

$$\mathcal{I}_P \subset \{\emptyset, \mathbb{U}_A\} \otimes \{\emptyset, \mathbb{U}_P\} \otimes \mathcal{F}$$

- As a consequence, the system is sequential
  - with the principal P as first player (leader)
  - and the agent A as second player (follower)
- Stackelberg games can be solved by bi-level optimization, for some information structures, like when

$$\mathcal{I}_P \vee \{\emptyset, \mathbb{U}_A\} \otimes \mathcal{U}_P \otimes \{\emptyset, \Omega\} \subset \mathcal{I}_A$$

Moral hazard

- An insurance company (the principal P) cannot observe the efforts of the insured (the agent A) to avoid risky behavior
- The firm faces the hazard that insured persons behave "immorally" (playing with matches at home)
- Moral hazard (hidden action) occurs when the decisions of the agent A are hidden to the principal P

$$\mathcal{I}_P \subset \{\emptyset, \mathbb{U}_A\} \otimes \{\emptyset, \mathbb{U}_P\} \otimes \mathcal{F}$$

- In case of moral hazard, the system is sequential with the principal as first player, (which does not preclude to choose the agent as first player in some special cases, as in a static team situation)
- Moral hazard games can be solved by bi-level optimization, for some information structures

Adverse selection

- In the absence of observable information on potential customers (the agent A), an insurance company (the principal P) offers a unique price for a contract hence screens and selects the "bad" ones
- Adverse selection occurs when
  - the agent A knows the state of nature (his type, or private information)
    $$\{\emptyset, \mathbb{U}_A\} \otimes \{\emptyset, \mathbb{U}_P\} \otimes \mathcal{F} \subset \mathcal{I}_A$$
    (the agent A can possibly observe the principal P action)
  - but the principal P does not know the state of nature
    $$\mathcal{I}_P \subset \mathcal{U}_A \otimes \{\emptyset, \mathbb{U}_P\} \otimes \{\emptyset, \Omega\}$$
    (the principal P can possibly observe the agent A action)
- In case of adverse selection, the system may or may not be sequential

Signaling

- In biology, a peacock signals its "good genes" (genotype) by its lavish tail (phenotype)
- In economics, a worker signals his working ability (productivity) by his educational level (diplomas)
- There is room for signaling
  - when the agent A knows the state of nature (his type)

$$\{\emptyset, \mathbb{U}_A\} \otimes \{\emptyset, \mathbb{U}_P\} \otimes \mathcal{F} \subset \mathcal{I}_A$$

    (the agent A can possibly observe the principal P action)
  - whereas the principal P does not know the state of nature, but the principal P observes the agent A action

$$\mathcal{I}_P = \mathcal{U}_A \otimes \{\emptyset, \mathbb{U}_P\} \otimes \{\emptyset, \Omega\}$$

as the agent A may reveal his type
by his decision which is observable by the principal P

Signaling

- The system is sequential (with the agent as first player) when

$$\mathcal{I}_A = \{\emptyset, \mathbb{U}_A\} \otimes \{\emptyset, \mathbb{U}_P\} \otimes \mathcal{F}$$

- The system is non causal when

$$\{\emptyset, \mathbb{U}_A\} \otimes \{\emptyset, \mathbb{U}_P\} \otimes \mathcal{F} \subsetneq \mathcal{I}_A \subset \{\emptyset, \mathbb{U}_A\} \otimes \mathcal{U}_P \otimes \mathcal{F}$$

Outline of the presentation

Stackelberg leadership model

- In the Stackelberg leadership model of game theory, we consider a leader P (principal) and a follower A (agent)
- We introduce the history space $\mathbb{H}$ equipped with the history field $\mathcal{H}$

$$\mathbb{H} = \mathbb{U}_A \times \mathbb{U}_P \times \Omega \; , \;\; \mathcal{H} = \mathcal{U}_A \otimes \mathcal{U}_P \otimes \mathcal{F}$$

- Denote by $\mathbf{U}_A$, $\mathbf{U}_P$ and $\mathbf{N}$ the coordinate mappings

$$\mathbf{U}_A : \mathbb{U}_A \times \mathbb{U}_P \times \Omega \to \mathbb{U}_A$$
$$\mathbf{U}_P : \mathbb{U}_A \times \mathbb{U}_P \times \Omega \to \mathbb{U}_P$$
$$\mathbf{N} : \mathbb{U}_A \times \mathbb{U}_P \times \Omega \to \Omega$$

- We exclude self-information, that is, we suppose that

$$\mathcal{I}_P \subset \sigma(\mathbf{U}_A, \mathbf{N}) \; , \;\; \mathcal{I}_A \subset \sigma(\mathbf{U}_P, \mathbf{N})$$

Stackelberg leadership model

- We suppose that $\langle P \rangle_{\mathfrak{P}} = \emptyset$,
  that is, leader P plays first,

$$\mathcal{I}_P \subset \{\emptyset, \mathbb{U}_A\} \otimes \{\emptyset, \mathbb{U}_P\} \otimes \mathcal{F}$$

  or, equivalently,

$$\mathcal{I}_P \subset \sigma(\mathbf{N})$$

- and that $\langle A \rangle_{\mathfrak{P}} = \{P\}$,
  that is, follower A plays second

$$\mathcal{I}_A \bigcap \{\emptyset, \mathbb{U}_A\} \otimes \mathcal{U}_P \otimes \{\emptyset, \Omega\} \neq \emptyset$$

  or, equivalently,

$$\mathcal{I}_A \cap \sigma(\mathbf{U}_P) \neq \emptyset \ , \ \ \mathcal{I}_A \subset \sigma(\mathbf{U}_P, \mathbf{N})$$

Admissible strategies and solution map

- We consider a couple $(\lambda_A, \lambda_P) \in \Lambda_A^{ad} \times \Lambda_P^{ad}$ of admissible strategies
- As $\mathcal{I}_A \subset \sigma(\mathbf{U}_P, \mathbf{N})$,
  the strategy λ_A of the follower A can be identified with

$$\lambda_A : \mathbb{U}_P \times \Omega \to \mathbb{U}_A$$

- We deduce that

$$\Omega \xrightarrow{S_{\lambda_P}} \mathbb{U}_P \times \Omega \xrightarrow{(\lambda_A, \mathrm{Id}_{\mathbb{U}_P \times \Omega})} \mathbb{U}_A \times \mathbb{U}_P \times \Omega$$

- The corresponding solution map satisfies

$$S_{(\lambda_A, \lambda_P)} = (\lambda_A \circ S_{\lambda_P}, S_{\lambda_P}) = (\lambda_A, \mathrm{Id}_{\mathbb{U}_P \times \Omega}) \circ S_{\lambda_P}$$

Strategy independence of conditional expectation (SICE)

We introduce the partial history space $\widetilde{\mathbb{H}}$
equipped with the partial history field $\widetilde{\mathcal{H}}$

$$\widetilde{\mathbb{H}} = \mathbb{U}_P \times \Omega \; , \;\; \widetilde{\mathcal{H}} = \mathcal{U}_P \otimes \mathcal{F}$$

and the partial information fields

$$\begin{aligned} \mathcal{I}_P = & \; \{\emptyset, \mathbb{U}_A\} \otimes \widetilde{\mathcal{I}_P} && \text{with } \widetilde{\mathcal{I}}_P \subset \{\emptyset, \mathbb{U}_P\} \otimes \mathcal{F} \subset \widetilde{\mathcal{H}} \\ \mathcal{I}_A = & \; \{\emptyset, \mathbb{U}_A\} \otimes \widetilde{\mathcal{I}}_A && \text{with } \widetilde{\mathcal{I}}_A \subset \mathcal{U}_P \otimes \mathcal{F} = \widetilde{\mathcal{H}} \end{aligned}$$

Assumption SICE

There exists a probability $\mathbb{Q}$ on $\widetilde{\mathbb{H}} = \mathbb{U}_P \times \Omega$ such that

$$\mathbb{P}_A \circ S_{\lambda_P}^{-1} = T_{\lambda_P}\mathbb{Q} \; , \;\; \mathbb{E}_{\mathbb{Q}}\big[T_{\lambda_P} \mid \widetilde{\mathcal{I}}_A\big] > 0 \; , \;\; \forall \lambda_P \in \Lambda_P^{ad}$$

and that

$$\mathbb{E}_{\mathbb{Q}}\big[j_A(u_A, \cdot) \mid \widetilde{\mathcal{I}}_A\big] = \mathbb{E}_{\mathbb{Q}}\big[j_A(u_A, \cdot) \mid \widetilde{\mathcal{I}}_A \vee \widetilde{\mathcal{I}}_P \vee \widetilde{\mathcal{D}}_P\big] \; , \;\; \forall u_A \in \mathbb{U}_A$$

Bayesian Nash equilibrium under assumption SICE

Bayesian Nash equilibrium

Under assumption SICE,
the couple $(\overline{\lambda}_A, \overline{\lambda}_P) \in \Lambda_A^{ad} \times \Lambda_P^{ad}$ of admissible strategies
is a Bayesian Nash equilibrium if (in case of payoffs)

$$\mathbb{E}_{\mathbb{Q}}\left[j_A \circ (\overline{\lambda}_A, \mathrm{Id}_{\mathbb{U}_P \times \Omega})\right] \geq \mathbb{E}_{\mathbb{Q}}\left[j_A \circ (\lambda_A, \mathrm{Id}_{\mathbb{U}_P \times \Omega})\right] \quad \forall \lambda_A \in \Lambda_A^{ad}$$

$$\mathbb{E}_{\mathbb{P}_P}\left[j_P \circ S_{(\overline{\lambda}_A, \overline{\lambda}_P)}\right] \geq \mathbb{E}_{\mathbb{P}_P}\left[j_P \circ S_{(\overline{\lambda}_A, \lambda_P)}\right] \quad \forall \lambda_P \in \Lambda_P^{ad}$$

Bayesian Nash equilibria can be obtained by bi-level optimization under assumption SICE

Suppose assumption SICE holds true

- The (upper level) optimization problem for the follower A

$$\min_{u_A \in \mathbb{U}_A} \mathbb{E}_{\mathbb{Q}}\left[j_A(u_A, \cdot) \mid \widetilde{\mathfrak{I}}_A \right]$$

provides (under technical assumptions, by a measurable selection theorem) an $\widetilde{\mathfrak{I}}_A$-measurable solution

$$\overline{\lambda}_A : \mathbb{U}_P \times \Omega \to \mathbb{U}_A \ , \ \ \sigma(\overline{\lambda}_A) \subset \widetilde{\mathfrak{I}}_A$$

- Then, the (lower level) optimization problem for the leader P is

$$\min_{\lambda_P \in \Lambda_P^{ad}} \mathbb{E}_{\mathbb{P}_P}\left[j_P \circ S_{(\overline{\lambda}_A, \lambda_P)} \right]$$

Thank you :-)

Evaluation of the criterion of the follower A (I)

We suppose that the follower A has a belief $\mathbb{P}_A$

Proposition

$$\mathbb{E}_{\mathbb{P}_A}\Big[j_A \circ S_{(\lambda_A,\lambda_P)} \mid S^{-1}_{(\lambda_A,\lambda_P)}(\mathfrak{I}_A)\Big] = \mathbb{E}_{\mathbb{P}_A \circ S^{-1}_{\lambda_P}}\Big[j_A \circ (\lambda_A, \mathrm{Id}_{\mathbb{U}_P\times\Omega}) \mid \widetilde{\mathfrak{I}}_A\Big]$$

Indeed,

$$\begin{aligned}
&\mathbb{E}_{\mathbb{P}_A}\Big[j_A \circ S_{(\lambda_A,\lambda_P)} \mid S^{-1}_{(\lambda_A,\lambda_P)}(\mathfrak{I}_A)\Big] \\
&= \mathbb{E}_{\mathbb{P}_A}\Big[j_A \circ (\lambda_A, \mathrm{Id}_{\mathbb{U}_P\times\Omega}) \circ S_{\lambda_P} \mid S^{-1}_{\lambda_P} \circ (\lambda_A, \mathrm{Id}_{\mathbb{U}_P\times\Omega})^{-1}(\mathfrak{I}_A)\Big] \\
&= \mathbb{E}_{\mathbb{P}_A \circ S^{-1}_{\lambda_P}}\Big[j_A \circ (\lambda_A, \mathrm{Id}_{\mathbb{U}_P\times\Omega}) \mid (\lambda_A, \mathrm{Id}_{\mathbb{U}_P\times\Omega})^{-1}(\mathfrak{I}_A)\Big]
\end{aligned}$$

Evaluation of the criterion of the follower A (II)

We have that

$$\mathcal{I}_A = \{\emptyset, \mathbb{U}_A\} \otimes \widetilde{\mathcal{I}}_A \text{ with } \widetilde{\mathcal{I}}_A \subset \mathcal{U}_P \otimes \mathcal{F}$$

and that

$$\big(\mathbb{U}_P \times \Omega, \mathcal{U}_P \otimes \mathcal{F}\big) \xrightarrow{(\lambda_A, \mathrm{Id}_{\mathbb{U}_P \times \Omega})} \big(\mathbb{U}_A \times \mathbb{U}_P \times \Omega, \mathcal{U}_A \otimes \mathcal{U}_P \otimes \mathcal{F}\big)$$

so that

$$(\lambda_A, \mathrm{Id}_{\mathbb{U}_P \times \Omega})^{-1}(\mathcal{I}_A) = \widetilde{\mathcal{I}}_A \vee \lambda_A^{-1}(\mathcal{I}_A) = \widetilde{\mathcal{I}}_A \vee \{\emptyset, \mathbb{U}_P\} \otimes \{\emptyset, \Omega\} = \widetilde{\mathcal{I}}_A$$

and we conclude that

$$\mathbb{E}_{\mathbb{P}_A}\Big[j_A \circ S_{(\lambda_A, \lambda_P)} \mid S^{-1}_{(\lambda_A, \lambda_P)}(\mathcal{I}_A)\Big] = \mathbb{E}_{\mathbb{P}_A \circ S^{-1}_{\lambda_P}}\Big[j_A \circ (\lambda_A, \mathrm{Id}_{\mathbb{U}_P \times \Omega}) \mid \widetilde{\mathcal{I}}_A\Big]$$

Evaluation of the criterion of the follower A (III)

Assumption

There exists a probability $\mathbb{Q}$ on $\widetilde{\mathbb{H}} = \mathbb{U}_P \times \Omega$ such that

$$\mathbb{P}_A \circ S_{\lambda_P}^{-1} = T_{\lambda_P}\mathbb{Q} \ , \ \ \mathbb{E}_{\mathbb{Q}}\big[T_{\lambda_P} \mid \widetilde{\mathfrak{I}}_A \big] > 0 \ , \ \ \forall \lambda_P \in \Lambda_P^{ad}$$

and that

$$\mathbb{E}_{\mathbb{Q}}\big[j_A(u_A, \cdot) \mid \widetilde{\mathfrak{I}}_A \big] = \mathbb{E}_{\mathbb{Q}}\big[j_A(u_A, \cdot) \mid \widetilde{\mathfrak{I}}_A \vee \widetilde{\mathfrak{I}}_P \vee \widetilde{\mathfrak{D}}_P \big] \ , \ \ \forall u_A \in \mathbb{U}_A$$

Proposition

$$\mathbb{E}_{\mathbb{P}_A \circ S_{\lambda_P}^{-1}}\Big[j_A \circ (\lambda_A, \mathrm{Id}_{\mathbb{U}_P \times \Omega}) \mid \widetilde{\mathfrak{I}}_A \Big] = \mathbb{E}_{\mathbb{Q}}\Big[j_A \circ (\lambda_A, \mathrm{Id}_{\mathbb{U}_P \times \Omega}) \mid \widetilde{\mathfrak{I}}_A \Big]$$

Evaluation of the criterion of the follower A (IV)

As $\mathbb{P}_A \circ S_{\lambda_P}^{-1} = T_{\lambda_P}\mathbb{Q}$ with $\mathbb{E}_{\mathbb{Q}}\big[T_{\lambda_P} \mid \widetilde{\mathfrak{I}}_A\big] > 0$, we have that

$$\mathbb{E}_{\mathbb{P}_A \circ S_{\lambda_P}^{-1}}\Big[j_A \circ (\lambda_A, \mathrm{Id}_{\mathbb{U}_P \times \Omega}) \mid \widetilde{\mathfrak{I}}_A\Big] = \frac{\mathbb{E}_{\mathbb{Q}}\Big[T_{\lambda_P} j_A \circ (\lambda_A, \mathrm{Id}_{\mathbb{U}_P \times \Omega}) \mid \widetilde{\mathfrak{I}}_A\Big]}{\mathbb{E}_{\mathbb{Q}}\Big[T_{\lambda_P} \mid \widetilde{\mathfrak{I}}_A\Big]}$$

Now, we also have

$$\begin{aligned}
&\mathbb{E}_{\mathbb{Q}}\Big[T_{\lambda_P} j_A \circ (\lambda_A, \mathrm{Id}_{\mathbb{U}_P \times \Omega}) \mid \widetilde{\mathfrak{I}}_A \vee \widetilde{\mathfrak{I}}_P \vee \widetilde{\mathfrak{D}}_P\Big] \\
&= T_{\lambda_P}\mathbb{E}_{\mathbb{Q}}\Big[j_A \circ (\lambda_A, \mathrm{Id}_{\mathbb{U}_P \times \Omega}) \mid \widetilde{\mathfrak{I}}_A \vee \widetilde{\mathfrak{I}}_P \vee \widetilde{\mathfrak{D}}_P\Big] \text{ as } \sigma(T_{\lambda_P}) \subset \widetilde{\mathfrak{I}}_P \vee \widetilde{\mathfrak{D}}_P \\
&= T_{\lambda_P}\mathbb{E}_{\mathbb{Q}}\Big[j_A \circ (\lambda_A, \mathrm{Id}_{\mathbb{U}_P \times \Omega}) \mid \widetilde{\mathfrak{I}}_A\Big] \text{ by assumption as } \sigma(\lambda_A) \subset \widetilde{\mathfrak{I}}_A
\end{aligned}$$

so that we deduce that

$$\mathbb{E}_{\mathbb{Q}}\Big[T_{\lambda_P} j_A \circ (\lambda_A, \mathrm{Id}_{\mathbb{U}_P \times \Omega}) \mid \widetilde{\mathfrak{I}}_A\Big] = \mathbb{E}_{\mathbb{Q}}\Big[T_{\lambda_P} \mid \widetilde{\mathfrak{I}}_A\Big]\mathbb{E}_{\mathbb{Q}}\Big[j_A \circ (\lambda_A, \mathrm{Id}_{\mathbb{U}_P \times \Omega}) \mid \widetilde{\mathfrak{I}}_A\Big]$$

and conclude

Evaluation of the criterion of the follower A (V)

We have proven that

Assumption

If there exists a probability $\mathbb{Q}$ on $\widetilde{\mathbb{H}} = \mathbb{U}_P \times \Omega$ such that

$$\mathbb{P}_A \circ S_{\lambda_P}^{-1} = T_{\lambda_P}\mathbb{Q} \; , \; \; \mathbb{E}_{\mathbb{Q}}\big[T_{\lambda_P} \mid \widetilde{\mathfrak{I}}_A \big] > 0 \; , \; \; \forall \lambda_P \in \Lambda_P^{ad}$$

and that

$$\mathbb{E}_{\mathbb{Q}}\big[j_A(u_A, \cdot) \mid \widetilde{\mathfrak{I}}_A \big] = \mathbb{E}_{\mathbb{Q}}\big[j_A(u_A, \cdot) \mid \widetilde{\mathfrak{I}}_A \vee \widetilde{\mathfrak{I}}_P \vee \widetilde{\mathfrak{D}}_P \big] \; , \; \; \forall u_A \in \mathbb{U}_A$$

then,

$$\mathbb{E}_{\mathbb{P}_A}\Big[j_A \circ S_{(\lambda_A, \overline{\lambda}_P)} \mid S_{(\lambda_A, \overline{\lambda}_P)}^{-1}(\mathfrak{I}_A) \Big] = \mathbb{E}_{\mathbb{Q}}\Big[j_A \circ (\lambda_A, \mathrm{Id}_{\mathbb{U}_P \times \Omega}) \mid \widetilde{\mathfrak{I}}_A \Big]$$

As a consequence, the best answer $\lambda_A^\star$ of the follower A does not depend on the leader P strategy $\overline{\lambda}_P$

Bayesian Nash equilibrium (with conditioning)

We suppose that follower A and the leader P each have belief $\mathbb{P}_A$ and $\mathbb{P}_P$

Bayesian Nash equilibrium (with conditioning)

The couple $(\overline{\lambda}_A, \overline{\lambda}_P) \in \Lambda_A^{ad} \times \Lambda_P^{ad}$ of admissible strategies
is a Bayesian Nash equilibrium (with conditioning) if (in case of payoffs)

$$\mathbb{E}_{\mathbb{Q}}\left[j_A \circ (\overline{\lambda}_A, \mathrm{Id}_{\mathbb{U}_P \times \Omega}) \mid \widetilde{\mathfrak{I}}_A\right] \geq \mathbb{E}_{\mathbb{Q}}\left[j_A \circ (\lambda_A, \mathrm{Id}_{\mathbb{U}_P \times \Omega}) \mid \widetilde{\mathfrak{I}}_A\right] \quad \forall \lambda_A \in \Lambda_A^{ad}$$

$$\mathbb{E}_{\mathbb{P}_P}\left[j_P \circ S_{(\overline{\lambda}_A, \overline{\lambda}_P)} \mid S^{-1}_{(\overline{\lambda}_A, \overline{\lambda}_P)}(\mathfrak{I}_P)\right] \geq \mathbb{E}_{\mathbb{P}_P}\left[j_P \circ S_{(\overline{\lambda}_A, \lambda_P)} \mid S^{-1}_{(\overline{\lambda}_A, \lambda_P)}(\mathfrak{I}_P)\right] \quad \forall \lambda_P \in \Lambda_P^{ad}$$