Game Theory with Information: Introducing the Witsenhausen Intrinsic Model

Michel De Lara and Benjamin Heymann
Cermics, École des Ponts ParisTech
France

École des Ponts ParisTech

March 15, 2017

Information plays a crucial role in competition

- Information — who knows what and when — plays a crucial role in competitive contexts
- Concealing, cheating, lying, deceiving are effective strategies

Our goals are to

1. introduce the notion of game in intrinsic form
2. contribute to the analysis of decentralized, non-cooperative decision settings
3. provide a (very) general mathematical language for game theory and mechanism design

What is a game in intrinsic form?

- Nature, the source of all randomness, or *states of Nature*
- Agents, who
  - hold *information*
  - make *decisions*, by means of *admissible strategies*, those fueled by information
- Players, who
  - hold *beliefs* about states of Nature
  - hold a *subset of agents* under their exclusive control (*team of executives*)
  - hold *objectives*, that they achieve by selecting proper admissible strategies for the agents under their control

In Witsenhausen's intrinsic form of a game, there is no tree structure (whereas Kuhn's extensive form of a game relies on a tree)

Outline of the presentation

Ingredients of Witsenhausen intrinsic model

Players and Nash equilibrium in Witsenhausen intrinsic model

Open questions (and research agenda)

Outline of the presentation

Outline of the presentation

Ingredients of Witsenhausen intrinsic model

- Agents and decisions, Nature, history
- Information fields and stochastic systems
- Strategies and admissible strategies
- Solvability and solution map

Players and Nash equilibrium in Witsenhausen intrinsic model

- Players in Witsenhausen intrinsic model
- Nash equilibrium in Witsenhausen intrinsic model

Open questions (and research agenda)

We will distinguish an individual from an agent

- An individual who makes a first, followed by a second decision, is represented by two agents (two decision makers)
- An individual who makes a sequence of decisions — one for each period $t = 0, 1, 2, \ldots, T-1$ — is represented by T agents, labelled $t = 0, 1, 2, \ldots, T-1$
- N individuals — each i of whom makes a sequence of decisions, one for each period $t = 0, 1, 2, \ldots, T_i - 1$ — is represented by $\prod_{i=1}^{N} T_i$ agents, labelled by

$$(i, t) \in \bigcup_{j=1}^{N} \{j\} \times \{0, 1, 2, \ldots, T_j - 1\}$$

Agents, decisions and decision space

- Let A be a finite set, whose elements are called agents (or decision-makers)
- Each agent $a \in A$ is supposed to make one decision $u_a \in \mathbb{U}_a$ where
  - the set $\mathbb{U}_a$ is the set of decisions for agent a
  - and is equipped with a σ-field $\mathcal{U}_a$
- We define the decision space as the product set

  $$\mathbb{U}_A = \prod_{b \in A} \mathbb{U}_b$$

  equipped with the product decision field

  $$\mathcal{U}_A = \bigotimes_{b \in A} \mathcal{U}_b$$

Examples

- $A = \{0, 1, 2, \ldots, T-1\}$ (T sequential decisions)
- $A = \{\mathtt{Pr}, \mathtt{Ag}\}$ (principal-agent models)

States of Nature and history space

- A state of Nature (or uncertainty, or scenario) is $\omega \in \Omega$ where
  - the set Ω is a measurable set, the sample space,
  - equipped with a σ-field $\mathcal{F}$
    (at this stage of the presentation, we do not need probability distribution, as we focus only on information)
- The history space is the product space

$$\mathbb{H} = \mathbb{U}_A \times \Omega = \prod_{b \in A} \mathbb{U}_b \times \Omega$$

equipped with the product history field

$$\mathcal{H} = \mathcal{U}_A \otimes \mathcal{F} = \bigotimes_{b \in A} \mathcal{U}_b \otimes \mathcal{F}$$

Examples

States of Nature Ω can include
types of players, randomness, stochastic processes

One agent, two possible decisions, two states of Nature

- Agents
$$A = \{a\}$$
- Decision set and field
$$\mathbb{U}_a = \{u_a^1, u_a^2\} \ , \ \ \mathcal{U}_a = \{\emptyset, \{u_a^1, u_a^2\}, \{u_a^1\}, \{u_a^2\}\}$$
- Sample space and field
$$\Omega = \{\omega^1, \omega^2\} \ , \ \ \mathcal{F} = \{\emptyset, \{\omega^1, \omega^2\}, \{\omega^1\}, \{\omega^2\}\}$$
- History space and field
$$\mathbb{H} = \mathbb{U}_a \times \Omega = \{u_a^1, u_a^2\} \times \{\omega^1, \omega^2\} \ , \ \ \mathcal{H} = 2^{\mathbb{H}}$$

Two agents, two possible decisions, two states of Nature

- Agents

$$A = \{a, b\}$$

- Decision sets and fields

$$\mathbb{U}_a = \{u_a^1, u_a^2\} \ , \ \ \mathcal{U}_a = \{\emptyset, \{u_a^1, u_a^2\}, \{u_a^1\}, \{u_a^2\}\}$$

  and

$$\mathbb{U}_b = \{u_b^1, u_b^2\} \ , \ \ \mathcal{U}_b = \{\emptyset, \{u_b^1, u_b^2\}, \{u_b^1\}, \{u_b^2\}\}$$

- Sample space and field

$$\Omega = \{\omega^1, \omega^2\} \ , \ \ \mathcal{F} = \{\emptyset, \{\omega^1, \omega^2\}, \{\omega^1\}, \{\omega^2\}\}$$

- History space and field

$$\mathbb{H} = \mathbb{U}_a \times \mathbb{U}_b \times \Omega = \{u_a^1, u_a^2\} \times \{u_b^1, u_b^2\} \times \{\omega^1, \omega^2\} \ , \ \ \mathcal{H} = 2^{\mathbb{H}}$$

Two *players*, T stages

- Agents

$$A = \{p, q\} \times \{0, 1, \ldots, T-1\}$$

- Decision sets and fields

$$\mathbb{U}_{(p,t)} = \mathbb{R}^{n_p} \; , \; \mathcal{U}_{(p,t)} = \mathcal{B}^{\mathrm{o}}_{\mathbb{R}^{n_p}} \; , \; \forall t = 0, 1, \ldots, T-1$$

and

$$\mathbb{U}_{(q,t)} = \mathbb{R}^{n_q} \; , \; \mathcal{U}_{(q,t)} = \mathcal{B}^{\mathrm{o}}_{\mathbb{R}^{n_q}} \; , \; \forall t = 0, 1, \ldots, T-1$$

- Sample space and field $(\Omega, \mathcal{F})$
- History space and field

$$\mathbb{H} = \prod_{t=0}^{T-1} \mathbb{U}_{(p,t)} \times \prod_{t=0}^{T-1} \mathbb{U}_{(q,t)} \times \Omega \; , \; \mathcal{H} = \bigotimes_{t=0}^{T-1} \mathcal{U}_{(p,t)} \otimes \bigotimes_{t=0}^{T-1} \mathcal{U}_{(q,t)} \otimes \mathcal{F}$$

Outline of the presentation

Information fields

- The information field of agent $a \in A$ is a σ-field

$$\mathcal{I}_a \subset \mathcal{H}$$

- In this representation, $\mathcal{I}_a$ is a subfield of the history field $\mathcal{H}$ which represents the information available to agent a when he makes a decision
- Therefore, the information of agent a may depend
  - on the states of Nature
  - and on other agents' decisions

One agent, two possible decisions, two states of Nature

- History space and field

$$\mathbb{H} = \mathbb{U}_a \times \Omega = \{u_a^1, u_a^2\} \times \{\omega^1, \omega^2\} \;, \;\; \mathcal{F} = 2^{\Omega} \;, \;\; \mathcal{H} = 2^{\mathbb{H}}$$

- Case where agent a knows nothing

$$\mathcal{I}_a = \{\emptyset, \mathbb{U}_a\} \otimes \{\emptyset, \Omega\} = \{\emptyset, \{u_a^1, u_a^2\}\} \otimes \{\emptyset, \{\omega^1, \omega^2\}\}$$

- Case where agent a knows the state of Nature

$$\begin{aligned}
\mathcal{I}_a =& \{\emptyset, \mathbb{U}_a\} \otimes \mathcal{F} \\
=& \{\emptyset, \mathbb{U}_a\} \otimes \{\emptyset, \{\omega^1, \omega^2\}, \{\omega^1\}, \{\omega^2\}\} \\
=& \underbrace{\{\emptyset, \{u_a^1, u_a^2\}\}}_{\text{undistinguishable}} \otimes \underbrace{\{\emptyset, \{\omega^1, \omega^2\}, \{\omega^1\}, \{\omega^2\}\}}_{\text{distinguishable}}
\end{aligned}$$

Two agents, two possible decisions, two states of Nature
Nested information fields

- History space and field

$$\mathbb{H} = \mathbb{U}_a \times \mathbb{U}_b \times \Omega = \{u_a^1, u_a^2\} \times \{u_b^1, u_b^2\} \times \{\omega^1, \omega^2\} \ , \ \ \mathcal{H} = 2^{\mathbb{H}}$$

- Agent a knows the state of Nature

$$\mathcal{I}_a = \{\emptyset, \mathbb{U}_a\} \times \{\emptyset, \mathbb{U}_b\} \otimes \{\emptyset, \{\omega^1, \omega^2\}, \{\omega^1\}, \{\omega^2\}\}$$

  and agent b knows the state of Nature and what agent a does

$$\mathcal{I}_b = \{\emptyset, \{u_a^1, u_a^2\}, \{u_a^1\}, \{u_a^2\}\} \times \{\emptyset, \mathbb{U}_b\} \otimes \{\emptyset, \{\omega^1, \omega^2\}, \{\omega^1\}, \{\omega^2\}\}$$

- In this example, information fields are nested

$$\mathcal{I}_a \subset \mathcal{I}_b$$

  meaning that agent b knows what agent a knows

Two agents, two decisions, two states of Nature
Non nested information fields

- History space and field

$$\mathbb{H} = \mathbb{U}_a \times \mathbb{U}_b \times \Omega = \{u_a^1, u_a^2\} \times \{u_b^1, u_b^2\} \times \{\omega^1, \omega^2\} \ , \ \ \mathcal{H} = 2^{\mathbb{H}}$$

- Agent a only knows the state of Nature

$$\mathcal{I}_a = \{\emptyset, \mathbb{U}_a\} \times \{\emptyset, \mathbb{U}_b\} \otimes \{\emptyset, \{\omega^1, \omega^2\}, \{\omega^1\}, \{\omega^2\}\}$$

  and agent b only knows what agent a does

$$\mathcal{I}_b = \{\emptyset, \{u_a^1, u_a^2\}, \{u_a^1\}, \{u_a^2\}\} \times \{\emptyset, \mathbb{U}_b\} \otimes \{\emptyset, \{\omega^1, \omega^2\}\}$$

- Information fields are not nested, $\mathcal{I}_a \not\subset \mathcal{I}_b$,
  as they cannot be compared by inclusion

Classical information patterns in game theory

Two agents: the principal $\mathtt{Pr}$ (leader) and the agent $\mathtt{Ag}$ (follower)

- *Moral hazard* (the insurance company cannot observe if the insured plays with matches at home)

$$\mathcal{I}_{\mathtt{Pr}} \subset \{\emptyset, \mathbb{U}_{\mathtt{Ag}}\} \otimes \{\emptyset, \mathbb{U}_{\mathtt{Pr}}\} \otimes \mathcal{F}$$

- *Stackelberg leadership model*

$$\mathcal{I}_{\mathtt{Ag}} \subset \{\emptyset, \mathbb{U}_{\mathtt{Ag}}\} \otimes \mathcal{U}_{\mathtt{Pr}} \otimes \mathcal{F} \ , \ \ \mathcal{I}_{\mathtt{Pr}} \subset \{\emptyset, \mathbb{U}_{\mathtt{Ag}}\} \otimes \{\emptyset, \mathbb{U}_{\mathtt{Pr}}\} \otimes \mathcal{F}$$

- *Adverse selection* (the insurance company cannot observe if the insured has good health)

$$\{\emptyset, \mathbb{U}_{\mathtt{Ag}}\} \otimes \{\emptyset, \mathbb{U}_{\mathtt{Pr}}\} \otimes \mathcal{F} \subset \mathcal{I}_{\mathtt{Ag}} \ , \ \ \mathcal{I}_{\mathtt{Pr}} \subset \mathcal{U}_{\mathtt{Ag}} \otimes \{\emptyset, \mathbb{U}_{\mathtt{Pr}}\} \otimes \{\emptyset, \Omega\}$$

- *Signaling* (the peacock's tail signals his good genes)

$$\{\emptyset, \mathbb{U}_{\mathtt{Ag}}\} \otimes \{\emptyset, \mathbb{U}_{\mathtt{Pr}}\} \otimes \mathcal{F} \subset \mathcal{I}_{\mathtt{Ag}} \ , \ \ \mathcal{I}_{\mathtt{Pr}} = \mathcal{U}_{\mathtt{Ag}} \otimes \{\emptyset, \mathbb{U}_{\mathtt{Pr}}\} \otimes \{\emptyset, \Omega\}$$

Two *players*, T stages

- Agents
$$A = \{p, q\} \times \{0, 1, \dots, T\}$$
- Information fields (at most, past decisions and state of Nature)
$$\mathfrak{I}_{(p,t)} \subset \bigotimes_{s=0}^{t-1} \mathfrak{U}_{(p,s)} \otimes \bigotimes_{s=t}^{T} \{\emptyset, \mathbb{U}_{(p,s)}\} \otimes \bigotimes_{s=0}^{t-1} \mathfrak{U}_{(q,s)} \otimes \bigotimes_{s=t}^{T} \{\emptyset, \mathbb{U}_{(q,s)}\} \otimes \mathfrak{F}$$
$$\mathfrak{I}_{(q,t)} \subset \bigotimes_{s=0}^{t-1} \mathfrak{U}_{(p,s)} \otimes \bigotimes_{s=t}^{T} \{\emptyset, \mathbb{U}_{(p,s)}\} \otimes \bigotimes_{s=0}^{t-1} \mathfrak{U}_{(q,s)} \otimes \bigotimes_{s=t}^{T} \{\emptyset, \mathbb{U}_{(q,s)}\} \otimes \mathfrak{F}$$

Stochastic system

Stochastic system

A stochastic system is a collection consisting of

- a finite set A of agents
- states of Nature $(\Omega, \mathcal{F})$
- decision sets, fields and information fields $\{\mathbb{U}_a, \mathcal{U}_a, \mathcal{I}_a\}_{a \in A}$

We will consider stochastic systems that display absence of self-information

Absence of self-information

A stochastic system displays absence of self-information when

$$\mathcal{I}_a \subset \mathcal{U}_{A \setminus \{a\}} \otimes \mathcal{F}$$

for any agent $a \in A$

- Absence of self-information means that the information of agent a may depend on the states of Nature and on all the other agents' decisions but not on his own decision
- Absence of self-information makes sense once we have distinguished an individual from an agent (else, it would lead to paradoxes)

What land have we covered?
What comes next?

- The stage is in place; so are the actors
  - Nature
  - agents
  - information
- How can actors play?
  - admissible strategies
  - solvability

Outline of the presentation

Ingredients of Witsenhausen intrinsic model
- Agents and decisions, Nature, history
- Information fields and stochastic systems
- Strategies and admissible strategies
- Solvability and solution map

Players and Nash equilibrium in Witsenhausen intrinsic model
- Players in Witsenhausen intrinsic model
- Nash equilibrium in Witsenhausen intrinsic model

Open questions (and research agenda)

Information fuels admissible strategies

A strategy (or policy, control law, control design) for agent a
is a measurable mapping

$$\lambda_a : (\mathbb{H}, \mathcal{H}) \to (\mathbb{U}_a, \mathcal{U}_a)$$

Admissible strategy

An admissible strategy for agent a is a mapping

$$\lambda_a : (\mathbb{H}, \mathcal{H}) \to (\mathbb{U}_a, \mathcal{U}_a)$$

which is measurable w.r.t. the information field $\mathcal{I}_a$ of agent a, that is,

$$\lambda_a^{-1}(\mathcal{U}_a) \subset \mathcal{I}_a$$

This condition expresses the property that
an admissible strategy for agent a
may only depend upon the information $\mathcal{I}_a$ available to him

Set of admissible strategies

We denote the set of admissible strategies of agent a by

$$\Lambda_a^{ad} = \left\{ \lambda_a : (\mathbb{H}, \mathcal{H}) \to (\mathbb{U}_a, \mathcal{U}_a) \;\middle|\; \lambda_a^{-1}(\mathcal{U}_a) \subset \mathcal{I}_a \right\}$$

and the set of admissible strategies of all agents is

$$\Lambda_A^{ad} = \prod_{a \in A} \Lambda_a^{ad}$$

Examples of admissible strategies

Consider a stochastic system with two agents a and b,
and suppose that σ-fields $\mathcal{U}_a$, $\mathcal{U}_b$ and $\mathcal{F}$ contain singletons

- Absence of self-information

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

Then, admissible strategies λ_a and λ_b have the form

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

- Sequential

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

Then, admissible strategies λ_a and λ_b have the form

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

Outline of the presentation

Solvability (I)

- In the Witsenhausen's intrinsic model,
  agents make decisions in an order which is not fixed in advance
- Briefly speaking, solvability is the property that,
  for each state of Nature, the agents' decisions
  are uniquely determined by their admissible strategies
- The solvability property is crucial to develop Witsenhausen's theory:
  without the solvability property,
  we would not be able to determine the agents decisions

Solvability (II)

The solvability problem consists in finding

- for any collection $\lambda = \{\lambda_a\}_{a \in A} \in \Lambda_A^{ad}$ of admissible policies
- for any state of Nature $\omega \in \Omega$
- decisions $u \in \mathbb{U}_A$ satisfying the implicit ("closed loop") equation

$$u = \lambda(u, \omega)$$

or, equivalently,

$$u_a = \lambda_a(\{u_b\}_{b \in A}, \omega) \ , \ \ \forall a \in A$$

Solvability property

A stochastic system displays the solvability property when

$$\forall \lambda \in \Lambda_A^{ad} \ , \ \ \forall \omega \in \Omega \ , \ \ \exists ! u \in \mathbb{U}_A \ , \ \ u = \lambda(u, \omega)$$

Solvability and information patterns

- Sequential

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

in which case

$$u_a = \lambda_a(\cancel{u_a}, \cancel{u_b}, \omega) = \widetilde{\lambda}_a(\omega) \ , \ \ u_b = \lambda_b(u_a, \cancel{u_b}, \omega) = \widetilde{\lambda}_b(u_a, \omega)$$

always displays a unique solution (u_a, u_b),
whatever $\omega \in \Omega$ and $\widetilde{\lambda}_a$ and $\widetilde{\lambda}_b$

- Deadlock

$$\mathcal{I}_a = \{\emptyset, \mathbb{U}_a\} \otimes \mathcal{U}_b \otimes \{\emptyset, \Omega\} \ , \ \ \mathcal{I}_b = \mathcal{U}_a \otimes \{\emptyset, \mathbb{U}_b\} \otimes \{\emptyset, \Omega\}$$

in which case

$$u_a = \widetilde{\lambda}_a(u_b) \ , \ \ u_b = \widetilde{\lambda}_b(u_a)$$

may display zero solutions, one solution or multiple solutions,
depending on the functional properties of $\widetilde{\lambda}_a$ and $\widetilde{\lambda}_b$

Solvability makes it possible to define a solution map

Solution map

Suppose that the solvability property holds true.
We define the solution map

$$S_\lambda : \Omega \to \mathbb{H} \; ,$$

that maps states of Nature towards histories, by

$$(u, \omega) = S_\lambda(\omega) \iff u = \lambda(u, \omega) \; , \; \forall (u, \omega) \in \mathbb{U}_A \times \Omega$$

We include the state of Nature ω in the image of $S_\lambda(\omega)$, so that we map the set Ω towards the history space $\mathbb{H}$, making it possible to interpret $S_\lambda(\omega)$ as a history driven by the admissible strategy λ
(in classical control theory, a state trajectory is produced by a policy)
For example, in the sequential case,

$$S_\lambda(\omega) = (\widetilde{\lambda}_a(\omega), \widetilde{\lambda}_b(\widetilde{\lambda}_a(\omega), \omega), \omega)$$

What land have we covered?
What comes next?

- The stage is in place; so are the actors
  - Nature
  - agents
  - information
- Actors know how they can play
  - admissible strategies
  - solvability
- In a non-cooperative context, we need
  - objectives
  - beliefs
  - a notion of equilibrium

Outline of the presentation

Outline of the presentation

Players hold teams of executive agents, objective functions and beliefs

- The set of players is denoted by P
- Every player $p \in P$ has
  - a team of executive agents
    $$A_p \subset A ,$$
    where $(A_p)_{p \in P}$ forms a partition of the set A of agents
  - a criterion (objective function)
    $$j_p : \mathbb{H} \to \mathbb{R} ,$$
    a measurable function over the history space $\mathbb{H}$
  - a belief
    $$\mathbb{P}_p : \mathcal{F} \to [0, 1] ,$$
    a probability distribution over the states of Nature $(\Omega, \mathcal{F})$

Example: two players, one agent per player

- Agents
$$A = \{a, b\}$$
- Players
$$p = \{a\} \ , \ \ A_p = \{a\} \ , \ \ q = \{b\} \ , \ \ A_q = \{b\}$$
- Criteria
$$j_{\{a\}}(u_a, u_b, \omega) \ , \ \ j_{\{b\}}(u_a, u_b, \omega)$$
- Beliefs $\mathbb{P}_{\{a\}}$ and $\mathbb{P}_{\{b\}}$ over $(\Omega, \mathcal{F})$

Example: two players, T stages

- Agents
$$A = \{p, q\} \times \{0, 1, \dots, T-1\}$$
- Players
$$P = \{p, q\}$$
$$A_p = \{p\} \times \{0, 1, \dots, T-1\} , \quad A_q = \{q\} \times \{0, 1, \dots, T-1\}$$
- Criteria
$$j_p(u_{(p,0)}, \dots, u_{(p,T-1)}, u_{(q,0)}, \dots, u_{(q,T-1)}, \omega) = \sum_{t=0}^{T-1} \mathrm{L}_{p,t}(u_{(p,t)}, u_{(q,t)}, \omega)$$
$$j_q(u_{(p,0)}, \dots, u_{(p,T-1)}, u_{(q,0)}, \dots, u_{(q,T-1)}, \omega) = \sum_{t=0}^{T-1} \mathrm{L}_{q,t}(u_{(p,t)}, u_{(q,t)}, \omega)$$
- Beliefs $\mathbb{P}_p$ and $\mathbb{P}_q$ over $(\Omega, \mathcal{F})$

How player p evaluates an admissible strategies profile λ

- Measurable solution map attached to $\lambda \in \Lambda_A^{ad}$ is

$$S_\lambda : \Omega \to \mathbb{H}$$

- Measurable criterion (costs or payoffs) is

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

- The composition of criteria with the solution map provides a random variable

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

- The random variable can be integrated w.r.t. the belief $\mathbb{P}_p$, yielding

$$\mathbb{E}_{\mathbb{P}_p}\left[j_p \circ S_\lambda\right] \in \mathbb{R}$$

where $\mathbb{E}_{\mathbb{P}_p}$ denotes the mathematical expectation w.r.t. the probability $\mathbb{P}_p$ on $(\Omega, \mathcal{F})$

Outline of the presentation

Pure (admissible) strategies profiles

- A pure (admissible) strategy for player p is an element of

$$\Lambda^{ad}_{A_p} = \prod_{a \in A_p} \Lambda^{ad}_a$$

- The set of pure (admissible) strategies for all players is

$$\prod_{p \in P} \Lambda^{ad}_{A_p} = \prod_{p \in P} \prod_{a \in A_p} \Lambda^{ad}_a = \prod_{a \in A} \Lambda^{ad}_a = \Lambda^{ad}_A$$

- An (admissible) strategies profile is

$$\lambda = (\lambda_p)_{p \in P} \in \prod_{p \in P} \Lambda^{ad}_{A_p}$$

- When we focus on player p, we write

$$\lambda = (\lambda_p, \lambda_{-p}) \in \Lambda^{ad}_{A_p} \times \underbrace{\prod_{p' \neq p} \Lambda^{ad}_{A_{p'}}}_{\Lambda^{ad}_{A_{-p}}}$$

Pure Bayesian Nash equilibrium

We say that the pure (admissible) strategies profile

$$\overline{\lambda} = (\overline{\lambda}_p)_{p \in P} \in \prod_{p \in P} \Lambda_{A_p}^{ad}$$

is a Bayesian Nash equilibrium if (in case of payoffs), for all $p \in P$,

$$\mathbb{E}_{\mathbb{P}_p}\left[j_p \circ S_{(\overline{\lambda}_p, \overline{\lambda}_{-p})}\right] \geq \mathbb{E}_{\mathbb{P}_p}\left[j_p \circ S_{(\lambda_p, \overline{\lambda}_{-p})}\right], \quad \forall \lambda_p \in \Lambda_{A_p}^{ad}$$

Mixed (admissible) strategies profiles

- A mixed (admissible) strategy (or randomized strategy) for player p is an element of
$$\Delta(\Lambda^{ad}_{A_p}) = \Delta(\prod_{a \in A_p} \Lambda^{ad}_a)$$
the set of probability distributions over the set of (admissible) strategies of his executives in A_p
- The definition of mixed strategies for player p reflects his ability to coordinate his team of executives in A_p
- By contrast, behavioral (admissible) strategies for player p are
$$\prod_{a \in A_p} \Delta(\Lambda^{ad}_a) \subset \Delta(\prod_{a \in A_p} \Lambda^{ad}_a)$$
and they do not require any correlating procedure

Mixed (admissible) strategies for players

- The set of mixed (admissible) strategies profiles is

$$\prod_{p \in P} \Delta(\Lambda^{ad}_{A_p}) = \prod_{p \in P} \Delta(\prod_{a \in A_p} \Lambda^{ad}_{a})$$

- A mixed (admissible) strategies profile is

$$\mu = (\mu_p)_{p \in P} \in \prod_{p \in P} \Delta(\Lambda^{ad}_{A_p})$$

- When we focus on player p, we write

$$\mu = (\mu_p, \mu_{-p}) \in \Delta(\Lambda^{ad}_{A_p}) \times \prod_{p' \neq p} \Delta(\Lambda^{ad}_{A_{p'}})$$

Mixed Bayesian Nash equilibrium

We say that the mixed (admissible) strategies profile

$$\overline{\mu} = (\overline{\mu}_p)_{p \in P} \in \prod_{p \in P} \Delta\big(\Lambda^{ad}_{A_p}\big)$$

is a Bayesian Nash equilibrium if (in case of payoffs), for all $p \in P$,

$$\int_{\Lambda^{ad}_p \times \Lambda^{ad}_{-p}} \overline{\mu}_p(d\lambda_p) \otimes \overline{\mu}_{-p}(d\lambda_{-p})\, \mathbb{E}_{\mathbb{P}_p}\Big[j_p \circ S_{(\lambda_p, \lambda_{-p})}\Big] \geq$$

$$\int_{\Lambda^{ad}_p \times \Lambda^{ad}_{-p}} \mu_p(d\lambda_p) \otimes \overline{\mu}_{-p}(d\lambda_{-p})\, \mathbb{E}_{\mathbb{P}_p}\Big[j_p \circ S_{(\lambda_p, \lambda_{-p})}\Big]\ , \ \forall \mu_p \in \Delta\big(\Lambda^{ad}_p\big)$$

Technical difficulties

- With which σ-algebra $\mathcal{M}$ can we equip the set of admissible strategies Λ_A^{ad}? So that we can consider $\Delta\left(\Lambda_A^{ad}\right)$, the set of probability distributions over Λ_A^{ad}
- Is the solution map

$$\Lambda_A^{ad} \times \Omega \to \mathbb{H}\,, \;\; (\lambda, \omega) \mapsto S_\lambda(\omega)$$

  measurable w.r.t. $\mathcal{M} \otimes \mathcal{F}$?
- Do we have to restrict to a subset of Λ_A^{ad}?

What land have we covered?
What comes next?

Witsenhausen intrinsic games cover

- deterministic games (with finite or measurable decision sets)
- deterministic dynamic games (finite span time)
- stochastic games
- stochastic dynamic games (finite span time)
- games in Kuhn extensive form (finite span time)

For games with enumerable or continuous span time,
the Witsenhausen intrinsic model has to be adapted

Outline of the presentation

Research questions

- How should we talk about games using WIM?
  - Can we extend the Bayesian Nash Equilibrium concept to general risk measures?
  - Can we re-organize the games bestiary using WIM?
  - How does the notion of subgame perfect Nash equilibrium translate within this framework?
- WIM: game theoretical results
  - What would a Nash theorem be in the WIM setting?
  - When do we have a generalized "backward induction" mechanism?
  - Under proper sufficient conditions on the information structure (extension of perfect recall), can we restrict the search among behavioral strategies instead of mixed strategies?
- Applications of WIM
  - What kind of applications do we target?
  - Can we use the WIM framework for mechanism design?

Nash Equilibrium with general risk measures

- We denote real-valued random variables on $(\Omega, \mathcal{F})$ by

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

- A risk measure $\mathbb{G}_p$ for the player p is a mapping

$$\mathbb{G}_p : \mathbb{L}(\Omega, \mathcal{F}) \to \mathbb{R} \cup \{+\infty\}$$

- We say that the players mixed (admissible) strategies profile

$$\overline{\mu} = (\overline{\mu}_p)_{p \in P} \in \prod_{p \in P} \Delta\big(\Lambda^{ad}_{A_p}\big)$$

is a Nash equilibrium if (in case of payoffs), for all $p \in P$,

$$\int_{\Lambda^{ad}_p \times \Lambda^{ad}_{-p}} \overline{\mu}_p(d\lambda_p) \otimes \overline{\mu}_{-p}(d\lambda_{-p}) \, \mathbb{G}_p\Big[j_p \circ S_{(\lambda_p, \lambda_{-p})}\Big] \geq$$

$$\int_{\Lambda^{ad}_p \times \Lambda^{ad}_{-p}} \mu_p(d\lambda_p) \otimes \overline{\mu}_{-p}(d\lambda_{-p}) \, \mathbb{G}_p\Big[j_p \circ S_{(\lambda_p, \lambda_{-p})}\Big] \ , \ \forall \mu_p \in \Delta\big(\Lambda^{ad}_p\big)$$

Can we re-organize the games bestiary using WIM?

With four relations between agents, introduced by Witsenhausen,

- Precedence relation $\mathfrak{P}$
- Subsystem relation $\mathfrak{S}$
- Information-memory relation $\mathfrak{M}$
- Decision-memory relation $\mathfrak{D}$

we can provide a typology of systems

- Static team
- Station
- Sequential systems
- Partially nested systems
- Quasiclassical systems
- Classical systems
- Hierarchical systems
- Parallel coordinated systems

Subgames and subgame perfect Nash equilibrium

- A subgame can be defined
  thanks to the notion of subsystem of agents in the WIM setting
- What are the conditions on a subsystem
  — w.r.t. players and their criteria —
  that make it possible to define a subgame?

We obtain a Nash theorem in the WIM setting

Theorem

Any finite, solvable, Witsenhausen game has a mixed Nash equilibrium

Proof

- The set of strategies is finite, as strategies map the finite history set towards finite decision sets
- To each strategy profile, we associate a payoff vector
- We thus obtain a matrix game and we can apply Nash theorem

Generalized existence result of Nash equilibria

By discretization

- Discretize decisions sets and sample space, and equip them with trace σ-fields
- Introduce discretized history set and σ-field

Current difficulties:

- How can we discretize information fields so that they are subfields of the trace history σ-field?
- With which topology can we equip the sample space Ω and the set Λ_A^{ad} of admissible strategies?
- Can we prove continuity for the solution map

$$\Lambda_A^{ad} \times \Omega \to \mathbb{H} \ , \ \ (\lambda, \omega) \mapsto S_\lambda(\omega) \quad ?$$

- Do we have to restrict to a subset of Λ_A^{ad} (like continuous admissible strategies)?

Generalized existence result of Nash equilibria

By best-reply set-valued mapping

- Define the best-reply set-valued mapping

$$\prod_{p \in P} \Delta(\Lambda^{ad}_{A_p}) \rightrightarrows \prod_{p \in P} \Delta(\Lambda^{ad}_{A_p})$$

Current difficulties:

- With which topology can we equip the sample space Ω and the set Λ^{ad}_{A} of admissible strategies?
- Can we prove continuity for the solution map

$$\Lambda^{ad}_{A} \times \Omega \to \mathbb{H}\,, \quad (\lambda, \omega) \mapsto S_\lambda(\omega) \quad ?$$

- Do we have to restrict to a subset of Λ^{ad}_{A} (like continuous admissible strategies)?
- What are the properties of the best-reply set-valued mapping? (measurability, convexity, continuity)?
- What are the proper fixed point theorems for set-valued mappings?

When do we have a generalized "backward induction" mechanism?

- Witsenhausen introduced the notion of strategy independence of conditional expectation (SICE)
- He showed that SICE was a key assumption for a generalized "backward induction" mechanism in stochastic optimal control
- Under assumption SICE, we provide sufficient conditions for a two players Bayesian Nash equilibrium to be obtained by bi-level optimization

Behavioral vs mixed strategies

- Mixed strategies profiles are

$$\prod_{p \in P} \Delta\Big(\prod_{a \in A_p} \Lambda_a^{ad} \Big)$$

  and reflect the synchronization of his agents by the player
- Behavioral strategies profiles are

$$\prod_{p \in P} \prod_{a \in A_p} \Delta(\Lambda_a^{ad})$$

  and they do not require any correlating procedure
- Under proper sufficient conditions on the information structure — generalizing perfect recall — we expect to prove that some games can be solved over the smaller set of behavioral strategies profiles instead of the large set of mixed strategies profiles

$$\underbrace{\prod_{p \in P} \prod_{a \in A_p} \Delta(\Lambda_a^{ad})}_{behavioral} \subset \underbrace{\prod_{p \in P} \Delta\Big(\prod_{a \in A_p} \Lambda_a^{ad} \Big)}_{mixed}$$

What kind of applications do we target?

- The WIM is of particular interest for non sequential games
- In particular we envision applications for networks, auctions and decentralized energy systems

Mechanism design presented in the intrinsic framework

- The designer (= principal) can extend the natural history set, by offering new decisions to every agent (messages)
- He is free to extend the information fields of the agents as he wishes
- He can partly shape the objective functions of the players

Conclusion

- a rich language
- a lot of open questions, and a lot of things not yet properly defined
- we are looking for feedback

Thank you :-)