Game Theory with Information: Introducing the Witsenhausen Intrinsic Model

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

July 6, 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, which includes both games in Kuhn's extensive form and Bayesian games
2. contribute to the analysis of decentralized, non-cooperative decision settings
3. provide a (very) general mathematical language for game theory and mechanism design

Outline of the presentation

Outline of the presentation

H. S. Witsenhausen. On information structures, feedback and causality. *SIAM J. Control*, 9(2):149–160, May 1971.

Sequentiality and perfect memory are tacit assumptions in control-oriented works on dynamic games

In control-oriented works on dynamic games (in particular, stochastic control problems) one usually finds a "dynamic equation" describing the evolution of a "state" in response to decision (control) variables of the players and to random variables. One also finds "output equations" which define output variables for a player as functions of the state, decision and random variables. Then the information structure is defined by allowing each decision variable to be any desired (measurable) function of the output variables generated for that player up to that time. Such a setup assumes that the time order in which the various decisions variables are selected is fixed in advance. It assumes that each player acts as if he had responsability only for one station. It assumes that this station has perfect memory.

Going beyond sequentiality and perfect memory

For large complex systems these tacit assumptions are unlikely to hold. (. . .) The order in which the various agents of the various organizations will have to act cannot always be predicted, and the information available to different agents, even of the same organization, may be noncomparable in the sense that, of two agents, neither one knows everything his colleague knows.

Kuhn's answer: games in extensive form

These difficulties in specifying the information structure of a game were faced and overcome in the early days of game theory

- Von Neumann and Morgenstern (1944)
  - fixed sequencing of decisions
  - variables range over finite sets
- Kuhn (1953)
  - removes the restriction of fixed sequencing of decisions
  - variables range over finite sets
- Aumann (1964)
  - fixed sequencing of decisions
  - variables range over measurable sets

Witsenhausen's answer: games as multiple feedback loops

> *The decision process is considered as a feedback loop and the game is characterized by its interaction with the policies of the agents, without prejudging questions of chronological order.*

In the Kuhn formulation,

> *the tree describing the game is an expression of the general solution of the closed loop relations. (These relations map information into decisions by the policies, and decisions into information by the rules of the game). For any combination of policies one can find the corresponding outcome by following the tree along selected branches, and this is an explicit procedure. Thus the difficulties that might arise in solving the loop have been eliminated by defining the game in terms of a general unique solution which must be found before the model can be set up.*

References

H. S. Witsenhausen. The intrinsic model for discrete stochastic control: Some open problems. In A. Bensoussan and J. L. Lions, editors, *Control Theory, Numerical Methods and Computer Systems Modelling*, *Lecture Notes in Economics and Mathematical Systems*, 107:322–335, Springer-Verlag, 1975.

H. S. Witsenhausen. On information structures, feedback and causality. *SIAM J. Control*, 9(2):149–160, May 1971.

H. S. Witsenhausen. On Policy Independence of Conditional Expectations. *Information and Control*, 28(1):65–75, 1975.

P. Carpentier, J.-P. Chancelier, G. Cohen, M. De Lara. Stochastic Multi-Stage Optimization. At the Crossroads between Discrete Time Stochastic Control and Stochastic Programming. Springer-Verlag, Berlin, 2015.

What is a game in normal form?

- Players, where each player
  - implements a *strategy*
  - holds an *objective*, which depends on his own strategy and on all the other players strategies

What is a Bayesian game?

- Nature, the source of all randomness, or *states of Nature*
- Players, where each player
  - holds a *belief* about states of Nature
  - implements a *strategy*
  - holds an *objective*, which depends on his own strategy, on all the other players strategies and on the state of Nature

What is a game in extensive form?

- Nature, the source of all randomness, or *states of Nature*
  - equipped with a *probability*
- Players, where each player
  - implements a *strategy*
  - holds an *objective*
- Nodes of a rooted tree, where each node
  - is the *executive* of a single player or of Nature
  - holds a *decision set* and a *transition mapping*
  - holds an *information set* (subset of nodes)
- Strategies thap map each information set onto decision set

Kuhn's extensive form of a game relies on a rooted tree

What is a game in intrinsic form?

- Nature, the source of all randomness, or *states of Nature*
- Players, where each player
  - holds a *belief* about states of Nature
  - implements a *strategy*
  - holds an *objective*
- Agents, where each agent
  - is the *executive* of a single player
  - holds a *decision set*
  - holds an *information field*, subfield of the *configuration set*
- Strategies, where the strategy of each agent
  - maps configuration set onto decision set
  - is measurable with respect to the agent information field

In Witsenhausen's intrinsic form of a game, there is no tree structure

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?

Outline of the presentation

Why the Witsenhausen intrinsic model?

Ingredients of the Witsenhausen intrinsic model (WIM)

Game theory in the Witsenhausen intrinsic model setting

Open questions (and research agenda)

Conclusion

Outline of the presentation

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 (or configuration space) is the product space

$$\mathbb{H} = \mathbb{U}_A \times \Omega = \prod_{a \in A} \mathbb{U}_a \times \Omega$$

equipped with the product history field

$$\mathcal{H} = \mathcal{U}_A \otimes \mathcal{F} = \bigotimes_{a \in A} \mathcal{U}_a \otimes \mathcal{F}$$

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

One *player*, T stages

- Agents
$$A = \{0, 1, \ldots, T-1\}$$
- Decision sets and fields
$$\mathbb{U}_t = \mathbb{R}^n \; , \;\; \mathcal{U}_t = \mathcal{B}^{\mathrm{o}}_{\mathbb{R}^n} \; , \;\; \forall t = 0, 1, \ldots, T-1$$
- Sample space and field $(\Omega, \mathcal{F})$
- History space and field
$$\mathbb{H} = \prod_{t=0}^{T-1} \mathbb{U}_t \times \Omega \; , \;\; \mathcal{H} = \bigotimes_{t=0}^{T-1} \mathcal{U}_t \otimes \mathcal{F}$$

Two *players*, T stages

- Agents
$$A = \{p, q\} \times \{0, 1, \dots, T-1\}$$
- Decision sets and fields
$$\mathbb{U}_{(p,t)} = \mathbb{R}^{n_p} \;, \; \mathcal{U}_{(p,t)} = \mathcal{B}^{o}_{\mathbb{R}^{n_p}} \;, \; \forall t = 0, 1, \dots, T-1$$
and
$$\mathbb{U}_{(q,t)} = \mathbb{R}^{n_q} \;, \; \mathcal{U}_{(q,t)} = \mathcal{B}^{o}_{\mathbb{R}^{n_q}} \;, \; \forall t = 0, 1, \dots, T-1$$
- 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

$$\mathcal{I}_a \subset \mathcal{H} = \bigotimes_{a \in A} \mathcal{U}_a \otimes \mathcal{F}$$

- 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

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

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

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

- 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

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

  and agent b only knows what agent a does

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

- 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)

- *Stackelberg leadership model*

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

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

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

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

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

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

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

One *player*, T stages
Non anticipativity

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

Two *players*, T stages
Non anticipativity

- Agents

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

- Information fields (at most, past decisions and state of Nature)

$$\mathcal{I}_{(p,t)} \subset \bigotimes_{s=0}^{t-1} \mathcal{U}_{(p,s)} \otimes \bigotimes_{s=t}^{T} \{\emptyset, \mathbb{U}_{(p,s)}\} \otimes \bigotimes_{s=0}^{t-1} \mathcal{U}_{(q,s)} \otimes \bigotimes_{s=t}^{T} \{\emptyset, \mathbb{U}_{(q,s)}\} \otimes \mathcal{F}$$

$$\mathcal{I}_{(q,t)} \subset \bigotimes_{s=0}^{t-1} \mathcal{U}_{(p,s)} \otimes \bigotimes_{s=t}^{T} \{\emptyset, \mathbb{U}_{(p,s)}\} \otimes \bigotimes_{s=0}^{t-1} \mathcal{U}_{(q,s)} \otimes \bigotimes_{s=t}^{T} \{\emptyset, \mathbb{U}_{(q,s)}\} \otimes \mathcal{F}$$

Action-information system (A-I system)

Stochastic system (Witsenhausen terminology)

An action-information system (A-I 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 A-I systems that display absence of self-information

Absence of self-information

An A-I system displays absence of self-information when

$$\mathcal{I}_a \subset \mathcal{U}_{A\setminus\{a\}} \otimes \mathcal{F} = \{\emptyset, \mathbb{U}_a\} \otimes \bigotimes_{b \in A\setminus\{a\}} \mathcal{U}_b \otimes \mathcal{F}$$

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)

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 $\mathtt{Pr}$ (leader), makes decisions $u_{\mathtt{Pr}} \in \mathbb{U}_{\mathtt{Pr}}$, where the set of decisions is equipped with a σ-field $\mathcal{U}_{\mathtt{Pr}}$
  - the agent $\mathtt{Ag}$ (follower) makes decisions $u_{\mathtt{Ag}} \in \mathbb{U}_{\mathtt{Ag}}$, where the set of decisions is equipped with a σ-field $\mathcal{U}_{\mathtt{Ag}}$
- and Nature, corresponding to private information (or type) of the agent $\mathtt{Ag}$
  - 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

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

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

- By these expressions of the information fields
  - $\mathcal{I}_{\mathtt{Pr}}$ of the principal $\mathtt{Pr}$ (leader)
  - $\mathcal{I}_{\mathtt{Ag}}$ of the agent $\mathtt{Ag}$ (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 $\mathtt{Ag}$ may partly observe the action of the leader $\mathtt{Pr}$

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

- whereas the leader $\mathtt{Pr}$ observes at most the state of Nature

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

- As a consequence, the system is sequential
  - with the principal $\mathtt{Pr}$ as first player (leader)
  - and the agent $\mathtt{Ag}$ as second player (follower)
- Stackelberg games can be solved by bi-level optimization, for some information structures, like when (work in progress)

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

Moral hazard

- An insurance company (the principal $\mathtt{Pr}$) cannot observe the efforts of the insured (the agent $\mathtt{Ag}$) 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 $\mathtt{Ag}$ are hidden to the principal $\mathtt{Pr}$

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

- 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 (work in progress)

Adverse selection

- In the absence of observable information on potential customers (the agent $\mathtt{Ag}$), an insurance company (the principal $\mathtt{Pr}$) offers a unique price for a contract hence screens and selects the "bad" ones
- Adverse selection occurs when
  - the agent $\mathtt{Ag}$ knows the state of nature (his type, or private information)

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

    (the agent $\mathtt{Ag}$ can possibly observe the principal $\mathtt{Pr}$ action)
  - but the principal $\mathtt{Pr}$ does not know the state of nature

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

    (the principal $\mathtt{Pr}$ can possibly observe the agent $\mathtt{Ag}$ 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 $\mathtt{Ag}$ knows the state of nature (his type)

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

  (the agent $\mathtt{Ag}$ can possibly observe the principal $\mathtt{Pr}$ action)
  - whereas the principal $\mathtt{Pr}$ does not know the state of nature, but the principal $\mathtt{Pr}$ observes the agent $\mathtt{Ag}$ action

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

as the agent $\mathtt{Ag}$ may reveal his type
by his decision which is observable by the principal $\mathtt{Pr}$

Signaling

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

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

- The system is non causal when

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

What land have we covered?
What comes next?

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

Outline of the presentation

Information fuels adapted strategies

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

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

Adapted strategy

An adapted 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 adapted strategy for agent a
may only depend upon the information $\mathcal{I}_a$ available to him

Set of adapted strategies

We denote the set of adapted strategies of agent a by

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

and the set of adapted strategies of all agents is

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

Examples of adapted strategies

Consider an A-I 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, adapted 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, adapted 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 adapted 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
- The solvability property is a playability property

Solvability (II)

The solvability problem consists in finding

- for any collection $\lambda = \{\lambda_a\}_{a \in A} \in \Lambda_A^{ad}$ of adapted 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,

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

Solvability property

An A-I system displays the solvability property when

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

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 adapted strategy λ
(in classical control theory, a state trajectory is produced by a policy)

In the sequential case, the solution map is given by iterated composition

- In the sequential case

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

- adapted strategies λ_a and λ_b have the form

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

- so that the solution map is

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

- because the system of equations $u = \lambda(u, \omega)$ here writes

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

Outline of the presentation

Causality

In a causal system, agents are ordered, one playing after the other with available information depending only on agents acting earlier, but the order may depend upon the history

We lay out mathematical ingredients to define causality: Orderings and partial orderings

- Let $\mathbb{O}$ denote the set of total orderings of agents in A, that is, injective mappings from $\{1, \ldots, A^\sharp\}$ to A, where $A^\sharp = \mathrm{card}(A)$
- For $k \in \{1, \ldots, A^\sharp\}$, let $\mathbb{O}_k$ denote the set of k-orderings, that is, injective mappings from $\{1, \ldots, k\}$ to A (thus $\mathbb{O} = \mathbb{O}_{A^\sharp}$)
- There is a natural restriction mapping $\psi_k : \mathbb{O} \to \mathbb{O}_k$, the restriction of any ordering of A to the domain set $\{1, \ldots, k\}$

We lay out mathematical ingredients to define causality: History-orderings

- Define a history-ordering as a mapping $\varphi : \mathbb{H} \to \mathbb{O}$ from histories towards orderings
- Along each history $h \in \mathbb{H}$, the agents are ordered by $\varphi(h) \in \mathbb{O}$
- With any $k \in \{1, \dots, A^\sharp\}$ and k-ordering $\rho_k \in \mathbb{O}_k$, we associate the set $\mathbb{H}^{\varphi}_{k,\rho_k}$ of histories that induce the same order than ρ_k for the agents having a rank smaller or equal to k, that is,

$$\mathbb{H}^{\varphi}_{k,\rho_k} = \{h \in \mathbb{H} \mid \psi_k(\varphi(h)) = \rho_k\}$$

Now, we define causality

Causality

A stochastic system is causal
if there exists (at least one) history-ordering φ from $\mathbb{H}$ towards $\mathbb{O}$,
with the property that for any $k \in \{1, \ldots, A^\sharp\}$ and $\rho_k \in \mathbb{O}_k$,
the set $\mathbb{H}^{\varphi}_{k,\rho_k}$ satisfies

$$\mathbb{H}^{\varphi}_{k,\rho_k} \cap G \in \mathcal{U}_{\{\rho_k(1),\ldots,\rho_k(k-1)\}} \otimes \mathcal{F} \ , \ \ \forall G \in \mathcal{I}_{\rho_k(k)}$$

- In other words, when the first k agents are known and ordered by $(\rho_k(1), \ldots, \rho_k(k))$, the information $\mathcal{I}_{\rho_k(k)}$ of the agent $\rho_k(k)$ with rank k depends at most on the decisions of agents with rank $< k$, that is, $\rho_k(1)$, $\ldots$, $\rho_k(k-1)$
- We say that a stochastic system is sequential if it is causal with a constant history-ordering

A causal but non sequential system

- We consider a set of agents $A = \{a, b\}$ with

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

- The agents' information fields are given by

$$\mathcal{I}_a = \sigma(\{u_a^1, u_a^2\} \times \{u_b^1, u_b^2\} \times \{\omega^2\}, \{u_a^1, u_a^2\} \times \{u_b^1\} \times \{\omega^1\})$$
$$\mathcal{I}_b = \sigma(\{u_a^1, u_a^2\} \times \{u_b^1, u_b^2\} \times \{\omega^1\}, \{u_a^1\} \times \{u_b^1, u_b^2\} \times \{\omega^2\})$$

- When the state of Nature is ω^2, agent a only sees ω^2, whereas agent b sees ω^2 and the decision of a: thus a acts first, then b
- The reverse holds true when the state of Nature is ω^1
- A non constant history-ordering mapping $\varphi : \mathbb{H} \to \{(a, b), (b, a)\}$ is defined by (for any couple (u_a, u_b))

$$\varphi\Big((u_a, u_b, \omega^2)\Big) = (a, b) \ \text{ and } \ \varphi\Big((u_a, u_b, \omega^1)\Big) = (b, a)$$

- The system is causal but not sequential

Causality implies solvability

Proposition

Causality implies (recursive) solvability with a measurable solution map

Kuhn's extensive form of a game encapsulates causality

Solvability does *not* imply causality

- Agents

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

- Decision sets

$$\mathbb{U}_a = \mathbb{U}_b = \mathbb{U}_c = \{0, 1\}$$

- Sample space (deterministic case)

$$\Omega = \{\omega\}$$

- Information fields

$$\begin{aligned} \mathcal{I}_a &= \sigma\big(u_b(1 - u_c)\big) \\ \mathcal{I}_b &= \sigma\big(u_c(1 - u_a)\big) \\ \mathcal{I}_c &= \sigma\big(u_a(1 - u_b)\big) \end{aligned}$$

- As no information field is trivial, the system is not causal

This non causal system is solvable

- Adapted strategies

$$\lambda_a(u_a, u_b, u_c) = \widetilde{\lambda}_a\big(u_b(1-u_c)\big)$$
$$\lambda_b(u_a, u_b, u_c) = \widetilde{\lambda}_b\big(u_c(1-u_a)\big)$$
$$\lambda_c(u_a, u_b, u_c) = \widetilde{\lambda}_c\big(u_a(1-u_b)\big)$$

where $\widetilde{\lambda} : \{0,1\} \to \{0,1\}$, hence

$$(\widetilde{\lambda}_a, \widetilde{\lambda}_b, \widetilde{\lambda}_c) \in \{\mathrm{Id}, 1-\mathrm{Id}\}^3$$

- Solution map

$$S_{(\mathrm{Id},\mathrm{Id},\mathrm{Id})} = (0,0,0)$$
$$S_{(1-\mathrm{Id},\mathrm{Id},\mathrm{Id})} = (1,0,1)$$
$$S_{(1-\mathrm{Id},1-\mathrm{Id},\mathrm{Id})} = (0,1,0)$$
$$S_{(1-\mathrm{Id},1-\mathrm{Id},1-\mathrm{Id})} = (1,1,1)$$

Don Juan wants to get married![1]

- Don Juan p is considering giving a phone call to his ex-lovers q, r, asking them if they want to marry him
- Don Juan selects one of his ex-lovers in the set $\{q, r\}$ and phones her
- If the answer to the first phone call is "yes",
  Don Juan marries the first called ex-lover
  (and decides not to give a second phone call)
- If the answer to the first phone call is "no",
  Don Juan makes a second phone call to the remaining ex-lover
- In that case, the remaining ex-lover answers "yes" or "no"

[1] Thanks to Miquel Oliu Barton

Don Juan wants to get married!

Agents and decisions

- Agents

$$A = \{ \overbrace{p_1, p_2}^{\text{Don Juan}} , \overbrace{q, r}^{\text{ex-lovers}} \}$$

because player Don Juan p makes two decisions,
hence has two executive agents p_1, p_2

- No Nature, but finite decisions sets

$$\mathbb{U}_{p_1} = \{q, r\} , \ \mathbb{U}_{p_2} = \{q, r, \partial\} , \ \mathbb{U}_q = \{Y, N\} , \ \mathbb{U}_r = \{Y, N\}$$

  - Agent p_1 selects an ex-lover in the set $\mathbb{U}_{p_1} = \{q, r\}$ and phones her
  - Agent p_2 either stops (decision ∂) or selects an ex-lover in $\{q, r\}$
  - Agents q, r either say "yes" or "no", hence select a decision in the set $\{Y, N\}$

- The finite decisions sets $\mathbb{U}_{p_1}, \mathbb{U}_{p_2}, \mathbb{U}_q, \mathbb{U}_r$ are equipped with the complete finite σ-fields $\mathcal{U}_{p_1}, \mathcal{U}_{p_2}, \mathcal{U}_q, \mathcal{U}_r$

Don Juan wants to get married!

Information structure: Don Juan

- When agent Don Juan p_1 makes the first phone call, he knows nothing

$$\mathcal{I}_{p_1} = \{\emptyset, \mathbb{U}_{p_1}\} \otimes \{\emptyset, \mathbb{U}_{p_2}\} \otimes \{\emptyset, \mathbb{U}_q\} \otimes \{\emptyset, \mathbb{U}_r\}$$

- The agent Don Juan p_2 remembers who Don Juan p_1 called first, and knows the answer

$$\mathcal{I}_{p_2} = \underbrace{\mathcal{U}_{p_1}}_{\text{remembering}} \otimes \{\emptyset, \mathbb{U}_{p_2}\} \otimes \underbrace{\mathcal{U}_q \otimes \mathcal{U}_r}_{\text{knowing the answer}}$$

Don Juan wants to get married!

Information structure: ex-lovers

- If ex-lover q receives a phone call from Don Juan, she does not know if she was called first or second, hence she cannot distinguish the elements in the set

$$\{\underbrace{(q,q),(q,r),(q,\partial),}_{\text{called first}} \quad \underbrace{(r,q)}_{\text{called second}}\}$$

so that her information field is

$$\mathcal{I}_q = \{\emptyset, \underbrace{\{(q,q),(q,r),(q,\partial),(r,q)\}}_{\text{called}}, \underbrace{\{(r,r),(r,\partial)\}}_{\text{not called}}, \mathbb{U}_{p_1} \times \mathbb{U}_{p_2}\} \otimes \mathcal{U}_q \otimes \mathcal{U}_r$$

- Conversely

$$\mathcal{I}_r = \{\emptyset, \{(r,r),(r,q),(r,\partial),(q,r)\}, \{(q,q),(q,\partial)\}, \mathbb{U}_{p_1} \times \mathbb{U}_{p_2}\} \otimes \mathcal{U}_q \otimes \mathcal{U}_r$$

Don Juan wants to get married!

A causal but non sequential system

If Don Juan p_1 calls ex-lover q first, the agents play in the following order

$$p_1 \to q \to p_2 \to r$$

and conversely

- History

$$\mathbb{H} = \mathbb{U}_{p_1} \times \mathbb{U}_{p_2} \times \mathbb{U}_q \times \mathbb{U}_r$$

- History partition

$$\mathbb{H}_q = \{q\} \times \mathbb{U}_{p_2} \times \mathbb{U}_q \times \mathbb{U}_r \ , \ \ \mathbb{H}_r = \{r\} \times \mathbb{U}_{p_2} \times \mathbb{U}_q \times \mathbb{U}_r$$

- A non constant history-ordering mapping is

$$\varphi : \mathbb{H} \to \{(p_1, q, p_2, r), (p_1, r, p_2, q)\}$$

  such that

$$\varphi_{|\mathbb{H}_q} \equiv (p_1, q, p_2, r) \ , \ \ \varphi_{|\mathbb{H}_r} \equiv (p_1, r, p_2, r)$$

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
  - adapted strategies
  - solvability
- In a non-cooperative context, we need
  - players endowed with
    - objectives
    - beliefs

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 (or avatars)

    $$A_p \subset A$$

    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

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

- Players

$$P = \{p, q\}$$

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

- Criteria

$$j_p(u_{(p,0)}, \ldots, u_{(p,T-1)}, u_{(q,0)}, \ldots, u_{(q,T-1)}, \omega) = \sum_{t=0}^{T-1} \mathrm{L}_{p,t}(u_{(p,t)}, u_{(q,t)}, \omega)$$

$$j_q(u_{(p,0)}, \ldots, u_{(p,T-1)}, u_{(q,0)}, \ldots, u_{(q,T-1)}, \omega) = \sum_{t=0}^{T-1} \mathrm{L}_{q,t}(u_{(p,t)}, u_{(q,t)}, \omega)$$

- Beliefs $\mathbb{P}_p$ and $\mathbb{P}_q$ over $(\Omega, \mathcal{F})$

Outline of the presentation

How player p evaluates an adapted 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

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

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

Pure (adapted) strategies profiles

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

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

- The set of pure (adapted) 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 (adapted) 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

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

We can now turn a game in intrinsic form into a game in normal form

Every player $p \in P$ has

- a strategy set made of pure (adapted) strategies

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

- a payoff from (adapted) strategies profiles to the real numbers

$$\lambda \in \prod_{p \in P} \Lambda^{ad}_{A_p} \mapsto \mathbb{E}_{\mathbb{P}_p}\left[j_p \circ S_\lambda\right] = \mathbb{E}_{\mathbb{P}_p \circ S_\lambda^{-1}}\left[j_p\right]$$

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)
- Bayesian 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?

Outline of the presentation

Pure Bayesian Nash equilibrium

We say that the pure (adapted) strategies profile

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

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

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

Mixed (adapted) strategies profiles (or selecting pure strategies randomly)

- A mixed (adapted) strategy (or randomized strategy) for player p is an element of
$$\Delta(\Lambda^{ad}_{A_p}) = \Delta\big(\prod_{a \in A_p} \Lambda^{ad}_a\big)$$
the set of probability distributions over the set of (adapted) 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 (adapted) strategies for player p are
$$\prod_{a \in A_p} \Delta(\Lambda^{ad}_a) \subset \Delta\big(\prod_{a \in A_p} \Lambda^{ad}_a\big)$$
and they do not require any correlating procedure

Mixed (adapted) strategies for players

- The set of mixed (adapted) 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 (adapted) 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

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

We can now define a mixed Bayesian Nash equilibrium

Mixed Bayesian Nash equilibrium

We say that the mixed (adapted) strategies profile

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

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

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

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

$$\forall \mu_p \in \Delta\left(\Lambda^{ad}_p\right)$$

Technical difficulties

- With which σ-algebra $\mathcal{M}$ can we equip the set Λ_A^{ad} of adapted strategies? So that we can consider and manipulate $\Delta\big(\Lambda_A^{ad}\big)$, 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 the set Λ_A^{ad} of adapted strategies?

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 do sets of admissible strategies, for the trace discretized information fields, behave when discretization is refined?
- 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

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

- 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

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

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?

Outline of the presentation

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

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

- For example, the worst-case risk measure is

$$\mathbb{G}\big[\mathbf{X}\big] = \inf_{\omega \in \Omega} \mathbf{X}(\omega)$$

Nash Equilibrium with general risk measures

We say that the players mixed (admissible) strategies profile

$$\overline{\mu} = (\overline{\mu}_p)_{p \in P} \in \prod_{p \in P} \Delta\left(\Lambda_{A_p}^{ad}\right)$$

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

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

Outline of the presentation

Can we re-organize the games bestiary using WIM?

H. S. Witsenhausen. The intrinsic model for discrete stochastic control: Some open problems. In A. Bensoussan and J. L. Lions, editors

With four relations between agents,

- 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, etc.

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?

A subsystem is a subset of agents closed w.r.t. information

We define the information $\mathcal{I}_B \subset \mathcal{H}$ of the subset $B \subset A$ of agents by

$$\mathcal{I}_B = \bigvee_{b \in B} \mathcal{I}_b$$

that is, the smallest σ-fields that contains all the σ-fields $\mathcal{I}_b$, for $b \in B$

Subsystem (Witsenhausen, 1975)

A nonempty subset B of agents in A is a subsystem if the information field $\mathcal{I}_B$ at most depends on the decisions of the agents in B, that is,

$$\mathcal{I}_B \subset \mathcal{U}_B \otimes \mathcal{F} = \bigotimes_{b \in B} \mathcal{U}_b \otimes \bigotimes_{c \notin B} \{\emptyset, \mathbb{U}_c\} \otimes \mathcal{F}$$

Thus, the information received by agents in B depends upon states of Nature and decisions of members of B only

Example: subsystems in stochastic control

In stochastic control, when past information accumulates in a filtration from initial time $t = 0$ to horizon $t = T$,

- agents $\{0, 1, \ldots, t\}$ up to time t form a subsystem, as they do not require decisions made by agents in $\{t+1, \ldots, T\}$ to make their own decisions
- whereas agents in $\{t+1, \ldots, T\}$ do *not* form a subsystem, as they *do* need decisions made by agents in $\{0, 1, \ldots, t\}$ to make their own decisions

Subsystems allow to decompose A-I systems

If the A-I system

$$A\ ,\ \ (\Omega,\mathcal{F})\ ,\ \ \{\mathbb{U}_a,\mathcal{U}_a,\mathcal{I}_a\}_{a\in A}$$

possesses a subsystem $B\subset A$ of agents,

- we can identify any information field $\mathcal{I}_a$, $a\in B$,

$$\mathcal{I}_a\subset\bigotimes_{b\in B}\mathcal{U}_b\otimes\bigotimes_{c\notin B}\{\emptyset,\mathbb{U}_c\}\otimes\mathcal{F}\ \text{ with }\ \mathcal{I}_a\subset\bigotimes_{b\in B}\mathcal{U}_b\otimes\mathcal{F}$$

- we can define two partial A-I systems

agents	Nature	decision and information
B	$(\Omega,\mathcal{F})$	$\{\mathbb{U}_b,\mathcal{U}_b,\mathcal{I}_b\}_{b\in B}$
$A\backslash B$	$\left(\prod_{b\in B}\mathbb{U}_b\times\Omega,\bigotimes_{b\in B}\mathcal{U}_b\otimes\mathcal{F}\right)$	$\{\mathbb{U}_a,\mathcal{U}_a,\mathcal{I}_a\}_{a\in A\backslash B}$

Subsystems allow to decompose admissible strategies

- We write, for any strategy $\lambda \in \Lambda_A$, $\lambda = (\lambda_B, \lambda_{A\setminus B})$, where

$$\lambda_B : \mathbb{U}_{A\setminus B} \times \mathbb{U}_B \times \Omega \to \mathbb{U}_B \ , \ \ \lambda_{A\setminus B} : \mathbb{U}_{A\setminus B} \times \mathbb{U}_B \times \Omega \to \mathbb{U}_{A\setminus B}$$

- For the two partial A-I systems, we denote
  - Λ_B and $\Lambda_{A\setminus B}$ the sets of strategies
  - Λ_A^{ad} and $\Lambda_{A\setminus B}^{ad}$ the sets of admissible strategies

We suppose that all σ-fields include singletons

Proposition

When $B \subset A$ is a subsystem,

- *in any admissible strategy $\lambda = (\lambda_B, \lambda_{A\setminus B})$,*
  the strategy λ_B can be identified with

$$\lambda_B : \mathbb{U}_B \times \Omega \to \mathbb{U}_B \ \text{ that is, } \ \Lambda_A^{ad} \subset \Lambda_B \times \Lambda_{A\setminus B}$$

- *the set Λ_A^{ad} of admissible strategies can be naturally decomposed as*

$$\Lambda_A^{ad} = \Lambda_B^{ad} \times \Lambda_{A\setminus B}^{ad}$$

that is, as admissible strategies on the two partial A-I systems

The solvability property is inherited by partial A-I systems

Proposition

Suppose that

- *the A-I system* $\{\mathbb{U}_a, \mathfrak{U}_a, \mathfrak{I}_a\}_{a \in A}$ *displays the solvability property*
- *the subset* $B \subset A$ *is a subsystem*

Then each of the two partial A-I systems, with agents B *and* $A \backslash B$*, also displays the solvability property*

The solvability property induces partial solution maps

Proposition (Existence of partial solution maps)

When $B \subset A$ is a subsystem, and the strategy $\lambda = (\lambda_B, \lambda_{A \setminus B})$ is admissible, the two partial solution maps

$$S_{\lambda_B} : \Omega \to \mathbb{U}_B \times \Omega \ \text{ and } \ S_{\lambda_{A\setminus B}} : \mathbb{U}_B \times \Omega \to \mathbb{U}_{A\setminus B} \times \mathbb{U}_B \times \Omega$$

are defined by the two partial solvability properties

$$\begin{aligned} u_B &= \lambda_B\big(u_B, \omega\big) && \iff u_B = S_{\lambda_B}(\omega) \\ u_{A\setminus B} &= \lambda_{A\setminus B}\big(u_B, u_{A\setminus B}, \omega\big) && \iff u_{A\setminus B} = S_{\lambda_{A\setminus B}}\big(u_B, \omega\big) \end{aligned}$$

Subsystem, solvability and co-cycle property

Proposition (Co-cycle property of the solution map)

Suppose that

- *the A-I system* $\{\mathbb{U}_a, \mathfrak{U}_a, \mathfrak{I}_a\}_{a \in A}$ *displays the solvability property*
- *the subset* $B \subset A$ *is a subsystem*
- *the strategy* $\lambda = (\lambda_B, \lambda_{A \setminus B})$ *is admissible*

The solution map $S_\lambda : \Omega \to \mathbb{U}_B \times \Omega$ *and the two partial solution maps*

$$S_{\lambda_B} : \Omega \to \mathbb{U}_B \times \Omega \text{ and } S_{\lambda_{A \setminus B}} : \mathbb{U}_B \times \Omega \to \mathbb{U}_{A \setminus B} \times \mathbb{U}_B \times \Omega$$

satisfy the following co-cycle property

$$S_\lambda = S_{(\lambda_B, \lambda_{A \setminus B})} = S_{\lambda_{A \setminus B}} \circ S_{\lambda_B}$$

$$S_{(\lambda_B, \lambda_{A \setminus B})} : \Omega \xrightarrow{S_{\lambda_B}} \mathbb{U}_B \times \Omega \xrightarrow{S_{\lambda_{A \setminus B}}} \mathbb{U}_{A \setminus B} \times \mathbb{U}_B \times \Omega$$

Outline of the presentation

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

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

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

H. S. Witsenhausen. On Policy Independence of Conditional Expectations. *Information and Control*, 28(1):65–75, 1975.

- 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
- He showed that conditions, on the information structure, generalizing perfect recall ensured SICE (at least in discrete settings)
- Under assumption SICE, we provide sufficient conditions for a two players Bayesian Nash equilibrium to be obtained by bi-level optimization (Work in progress...)

Optimization problem

We suppose given a measurable criterion (objective function)

$$j : \mathbb{U}_A \times \Omega \to \mathbb{R}$$

We consider the optimization problem

$$\min_{\lambda_A \in \Lambda_A^{ad}} \mathbb{E}_{\mathbb{P}} \left[j \circ S_\lambda \right]$$

Strategy independence of conditional expectation (SICE)

Assumption SICE

1. There exists a probability $\mathbb{Q}_B$ on $\mathbb{U}_B \times \Omega$ such that

$$\mathbb{P} \circ S_{\lambda_B}^{-1} = T_{\lambda_B} \mathbb{Q}_B \text{ with } \mathbb{E}_{\mathbb{Q}_B}\big[T_{\lambda_B} \mid \mathcal{I}_B\big] > 0 \,, \;\; \forall \lambda_B \in \Lambda_B^{ad}$$

2. There exists a probability $\mathbb{Q}_A$ on $\mathbb{U}_A \times \Omega = \mathbb{U}_{A \setminus B} \times \mathbb{U}_B \times \Omega$ such that

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

In the discrete case, Witsenhausen provides sufficient conditions, on the information structure, to obtain SICE
H. S. Witsenhausen. On Policy Independence of Conditional Expectations. *Information and Control*, 28(1):65–75, 1975.

Dynamic programming equation (Work in progress...)

$$V_A = \mathbb{E}_{\mathbb{Q}_A}\left[j \mid \mathcal{I}_A\right]$$

$$V_B = \min_{\lambda_{A\setminus B} \in \Lambda^{ad}_{A\setminus B}} \mathbb{E}_{\mathbb{Q}_B}\left[V_A \circ S_{\lambda_{A\setminus B}} \mid \mathcal{I}_B\right]$$

$$V_\emptyset = \min_{\lambda_B \in \Lambda^{ad}_B} \mathbb{E}_{\mathbb{P}}\left[V_B \circ S_{\lambda_B}\right]$$

Outline of the presentation

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 :-)