Causal Inference Theory with Information Algebras: Introducing the Witsenhausen Intrinsic Model

Benjamin Heymann, Michel De Lara, Jean-Philippe Chancelier

BH: Criteo AI Lab, Paris, France
MDL and JPC : Cermics, École des Ponts, Marne-la-Vallée, France

Outline of the presentation

Witsenhausen intrinsic model [15']

Causality

Witsenhausen intrinsic model [15']

Witsenhausen intrinsic model [15’]

Agents, actions, Nature, configuration space, information σ-algebras

Agents, action spaces and Nature space

- Let A be a (finite or infinite) set, whose elements are called agents (or decision-makers)
- With each agent $a \in A$ is associated a measurable space

$$\big(\underbrace{\mathbb{U}_a}_{\substack{\text{set of}\\ \text{actions}\\ \text{for agent } a}} , \underbrace{\mathcal{U}_a}_{\substack{\sigma\text{-algebra}\\ \subset 2^{\mathbb{U}_a}}} \big)$$

- With Nature is associated a measurable space

$$(\Omega, \mathcal{F})$$

(at this stage of the presentation, we do not need to equip $(\Omega, \mathcal{F})$ with a probability distribution, as we only focus on information)

The configuration space is a product space

Configuration space

The configuration space is the product space

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

equipped with the product σ-algebra, called configuration σ-algebra

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

so that $(\mathbb{H}, \mathcal{H})$ is a measurable space

Example of configuration space

$\mathbb{U}_a = \{T_a, B_a\}$, $\mathbb{U}_b = \{R_b, L_b\}$, $\Omega = \{\omega^+, \omega^-\}$
$\mathcal{U}_a = 2^{\mathbb{U}_a}$, $\mathcal{U}_b = 2^{\mathbb{U}_b}$, $\mathcal{F} = 2^{\Omega}$

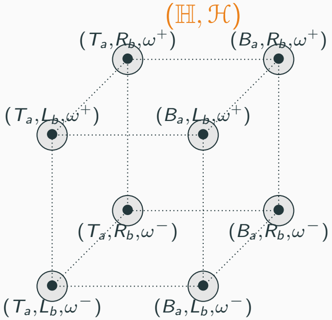

- product configuration space

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

- product configuration σ-algebra

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

represented by
the partition of its atoms

Information σ-algebras

Information σ-algebras express dependencies

Information σ-algebra of an agent

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

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

which is a sub σ-algebra of the product configuration σ-algebra

- The sub σ-algebra $\mathcal{I}_a$ of the configuration σ-algebra $\mathcal{H}$ represents the information available to agent a when the agent chooses an action
- Therefore, the information of agent a may depend
  - on the states of Nature
  - and on other agents' actions

In the finite case, information σ-algebras are represented by the partition of its atoms

The information σ-algebra of agent $a \in A$
is a sub σ-algebra $\mathcal{I}_a \subset \mathcal{H} = \bigotimes_{a \in A} \mathcal{U}_a \otimes \mathcal{F}$
which can, in the finite case, be represented by the partition of its atoms

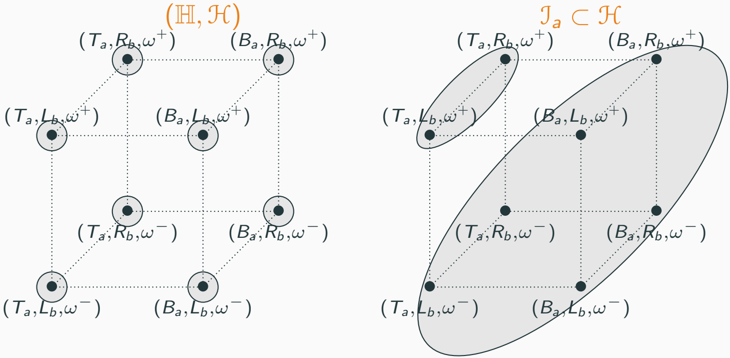

Elements of an atom cannot be distinguished by the agent a

Definition of the W-model (2 basic objects, possibly 1 axiom)

W-model

A W-model $\big(A, (\Omega, \mathcal{F}), (\mathbb{U}_a, \mathcal{U}_a)_{a \in A}, (\mathcal{I}_a)_{a \in A}\big)$
consists of 2 basic objects

(W-BO1a) the sample space $(\Omega, \mathcal{F})$

(W-BO1b) the collection $(\mathbb{U}_a, \mathcal{U}_a)_{a \in A}$ of agents' action spaces

(W-BO2) the collection $(\mathcal{I}_a)_{a \in A}$ of agents' information sub σ-algebras of $\mathcal{H} = \bigotimes_{a \in A} \mathcal{U}_a \otimes \mathcal{F}$

and (possibly) 1 axiom imposed on them

(W-Axiom1) for all agent $a \in A$, absence of self-information holds

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

To avoid paradoxes, we can consider W-models that display absence of self-information

Absence of self-information

A W-model displays absence of self-information when

$$\mathfrak{I}_a \subset \underbrace{\{\emptyset, \mathbb{U}_a\}}_{\text{not one's own action}} \otimes \underbrace{\bigotimes_{b \in A \setminus \{a\}} \mathfrak{U}_b}_{\text{other agents' actions}} \otimes \mathfrak{F} \,, \quad \forall a \in A$$

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

SIAM J. CONTROL
Vol. 9, No. 2, May 1971

ON INFORMATION STRUCTURES, FEEDBACK AND CAUSALITY*

H. S. WITSENHAUSEN†

Abstract. A finite number of decisions, indexed by $\alpha \in A$, are to be taken. Each decision amour to selecting a point in a measurable space $(U_\alpha, \mathscr{F}_\alpha)$. Each decision is based on some information f back from the system and characterized by a subfield $\mathscr{I}_\alpha$ of the product space $(\prod_\alpha U_\alpha, \prod_\alpha \mathscr{F}_\alpha)$. T decision function for each α can be any function γ_α measurable from $\mathscr{I}_\alpha$ to $\mathscr{F}_\alpha$.

A property of the $\{\mathscr{I}_\alpha\}_{\alpha \in A}$ is defined which assures that the setup has a causal interpretatio This property implies that for any combination of choices of the γ_α, the closed loop equations have unique solution.

The converse implication is false, when card $A > 2$.

1. Introduction. In control-oriented works on dynamic games (in particula stochastic control problems) one usually finds a "dynamic equation" describi the evolution of a "state" in response to decision (control) variables of the playe and to random variables. One also finds "output equations" which define outp variables for a player as functions of the state, decision and random variabl Then the information structure is defined by allowing each decision variable to any desired (measurable) function of the output variables generated for that play

Witsenhausen intrinsic model [15’]

Examples

Alice and Bob

"Alice and Bob" configuration space

Example

- no Nature
- two agents a (Alice) and b (Bob)
- two possible actions each $\mathbb{U}_a = \{T_a, B_a\}$, $\mathbb{U}_b = \{R_b, L_b\}$
- product configuration space (4 elements)

$$\mathbb{H} = \{T_a, B_a\} \times \{R_b, L_b\}$$

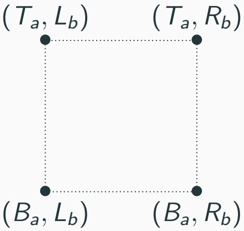

"Alice and Bob" information partitions

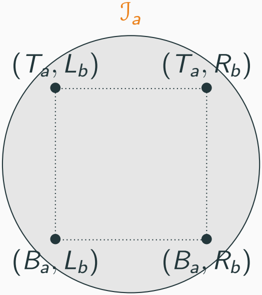

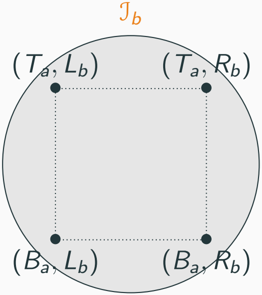

- $\mathcal{I}_a = \{\emptyset, \{T_a, B_a\}\} \otimes \{\emptyset, \{R_b, L_b\}\}$
  Alice knows nothing
- $\mathcal{I}_b = \{\emptyset, \{T_a, B_a\}\} \otimes \{\emptyset, \{R_b, L_b\}\}$
  Bob knows nothing

Alice knows Bob's action

"Alice and Bob" information partitions

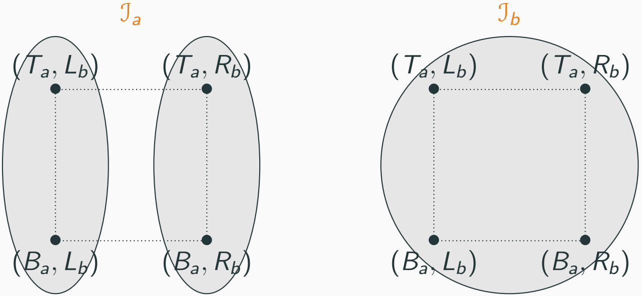

- $\mathcal{I}_b = \{\emptyset, \{T_a, B_a\}\} \otimes \{\emptyset, \{R_b, L_b\}\}$
  Bob knows nothing
- $\mathcal{I}_a = \{\emptyset, \{T_a, B_a\}\} \otimes \{\emptyset, \{R_b\}, \{L_b\}, \{R_b, L_b\}\}$
  Alice knows what Bob does
  (as she can distinguish between Bob's actions $\{R_b\}$ and $\{L_b\}$)

Alice, Bob and a coin tossing

"Alice, Bob and a coin tossing" configuration space

Example

- two states of Nature $\Omega = \{\omega^+, \omega^-\}$ (heads/tails)
- two agents a and b
- two possible actions each: $\mathbb{U}_a = \{T_a, B_a\}$, $\mathbb{U}_b = \{R_b, L_b\}$
- product configuration space (8 elements)

$$\mathbb{H} = \{T_a, B_a\} \times \{R_b, L_b\} \times \{\omega^+, \omega^-\}$$

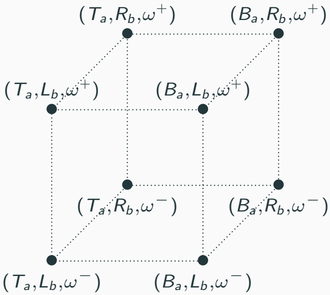

"Alice, Bob and a coin tossing" information partitions

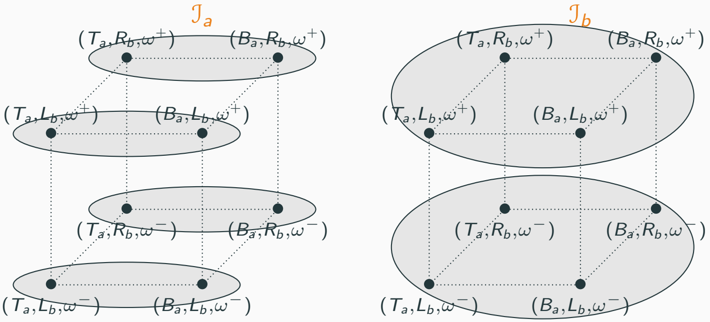

$$\mathfrak{I}_b = \overbrace{\{\emptyset, \{T_a, B_a\}\}}^{\text{Bob does not know what Alice does}} \otimes \{\emptyset, \mathbb{U}_b\} \otimes \overbrace{\{\emptyset, \{\omega^+\}, \{\omega^-\}, \{\omega^+, \omega^-\}\}}^{\text{Bob knows Nature's move}}$$

$$\mathfrak{I}_a = \{\emptyset, \mathbb{U}_a\} \otimes \underbrace{\{\emptyset, \{R_b\}, \{L_b\}, \{R_b, L_b\}\}}_{\text{Alice knows what Bob does}} \otimes \underbrace{\{\emptyset, \{\omega^+\}, \{\omega^-\}, \{\omega^+, \omega^-\}\}}_{\text{Alice knows Nature's move}}$$

Absent-minded driver

Absent-minded driver

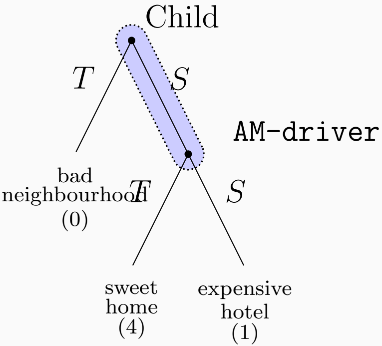

- S=Stay, T=Turn
- "paradox" that raised a problem in game theory
- the player looses public time, as plays "SS" "ST" cross the information set twice
- cannot be modelled *per se* in tree models (violates "no-AM" axiom)

A W-model for the absent-minded driver

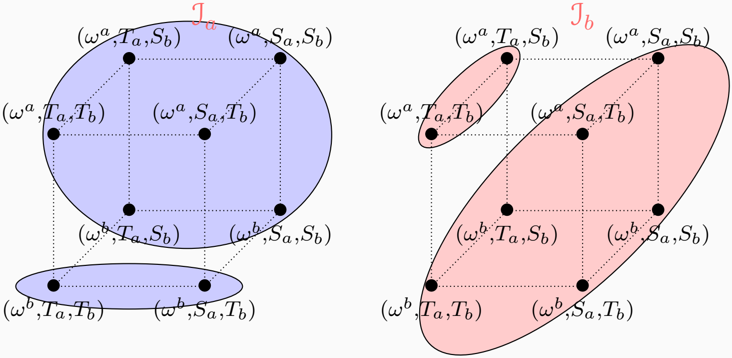

$$\mathcal{I}_a = \{\emptyset, \overbrace{\underbrace{\mathbb{U}_a \times \mathbb{U}_b \times \{\omega_a\}}_{\text{agent } a \text{ is whether the first one to act}} \cup \underbrace{\{S_b\} \times \mathbb{U}_a \times \{\omega_b\}}_{\text{or he acts second after agent } b \text{ has chosen } S}}^{\text{agent } a \text{ makes a move}}, \overbrace{\underbrace{\{T_b\} \times \mathbb{U}_a \times \{\omega_b\}}_{\text{agent } b \text{ chose } T \text{ and finished the game}}}^{\text{agent } a \text{ doesn't make a move}}, \mathbb{H}\}$$

$$\mathcal{I}_b = \{\emptyset, \mathbb{U}_a \times \mathbb{U}_b \times \{\omega_b\} \cup \{S_a\} \times \mathbb{U}_b \times \{\omega_a\}, \{T_a\} \times \mathbb{U}_b \times \{\omega_a\}, \mathbb{H}\}$$

What land have we covered?
What comes next?

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

Witsenhausen intrinsic model [15']

Strategies and solvability property

Information is the fuel of W-strategies

W-strategy of an agent

A (pure) W-strategy of 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 σ-algebra $\mathcal{I}_a$, that is,

$$\underbrace{\lambda_a^{-1}(\mathcal{U}_a)}_{\substack{\sigma\text{-algebra} \\ \text{generated by} \\ \text{W-strategy } \lambda_a}} \subset \underbrace{\mathcal{I}_a}_{\substack{\text{information} \\ \sigma\text{-algebra} \\ \text{of agent } a}}$$

This condition expresses the property that
a W-strategy $\lambda_a : (\mathbb{H}, \mathcal{H}) \to (\mathbb{U}_a, \mathcal{U}_a)$ for agent a
can only depend on the information $\mathcal{I}_a$ available to the agent

For instance, $\lambda_a^{-1}(\mathcal{U}_a) \subset \underbrace{\{\emptyset, \mathbb{H}\}}_{\text{no information}} \iff \lambda_a$ is constant on $\mathbb{H}$

Examples of W-strategies

Consider a W-model with two agents a and b,
and suppose that the σ-algebras $\mathcal{U}_a$, $\mathcal{U}_b$ and $\mathcal{F}$ contain the 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, W-strategies λ_a and λ_b have the form

$$\lambda_a(\not{u_a}, u_b, \omega) = \widetilde{\lambda}_a(u_b, \omega) \ , \ \ \lambda_b(u_a, \not{u_b}, \omega) = \widetilde{\lambda_b}(u_a, \omega)$$

- Sequential W-model

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

  Then, W-strategies λ_a and λ_b have the form

$$\lambda_a(\not{u_a}, u_b, \omega) = \widetilde{\lambda}_a(u_b, \omega) \ , \ \ \lambda_b(\not{u_a}, \not{u_b}, \omega) = \widetilde{\lambda}_b(\omega)$$

Set of W-strategies of an agent

We denote the set of (pure) W-strategies of agent a by

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

and the set of W-strategies of all agents is

$$\Lambda = \Lambda_A = \prod_{a \in A} \Lambda_a$$

Structural causal and Witsenhausen intrinsic models

Structural causal model	*Witsenhausen intrinsic model*
exogeneous variables	Nature $\omega \in \Omega$ (meas. space $(\Omega, \mathcal{F})$)
exogeneous distribution	
index of endogeneous variables	agent $a \in A$
domain of endogeneous variables	action set $\mathbb{U}_a$ (meas. space $(\mathbb{U}_a, \mathcal{U}_a)$)
	configuration space
	$\mathbb{H} = \prod_{a \in A} \mathbb{U}_a \times \Omega$, $\mathcal{H} = \bigotimes_{a \in A} \mathcal{U}_a \otimes \mathcal{F}$
	information σ-algebras $\{\mathcal{I}_a\}_{a \in A} \subset \mathcal{H}$
functional relation	W-strategy $\lambda_a : (\mathbb{H}, \mathcal{H}) \to (\mathbb{U}_a, \mathcal{U}_a)$
	$\lambda_a^{-1}(\mathcal{U}_a) \subset \mathcal{I}_a$, $\forall a \in A$
causal mechanism	W-strategy profile $\{\lambda_a\}_{a \in A}$

Solvability

Solvability

- In the Witsenhausen's intrinsic model,
  agents make actions in an order which is not fixed in advance
- Briefly speaking, solvability
  is the property that, for each state of Nature,
  the agents' actions are uniquely determined by their W-strategies

Solvability problem

The solvability problem consists in finding

- for any collection $\lambda = \{\lambda_a\}_{a \in A} \in \Lambda_A$ of W-strategies
- for any state of Nature $\omega \in \Omega$

actions $u \in \mathbb{U}_A$ satisfying
the implicit ("closed loop") equation

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

or, equivalently, the family of "closed loop" equations

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

Solvability property

Solvability property

A W-model displays the solvability property when

$$\forall \lambda = (\lambda_a)_{a \in A} \in \Lambda_A \;, \;\; \forall \omega \in \Omega \;, \;\; \exists ! u \in \mathbb{U}_A \;, \;\; u = \lambda(u, \omega)$$

or, equivalently,

$$\forall \lambda = (\lambda_a)_{a \in A} \in \Lambda_A \;, \;\; \forall \omega \in \Omega \;, \;\; \exists ! u \in \mathbb{U}_A$$

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

Solvability is a property of the information structure

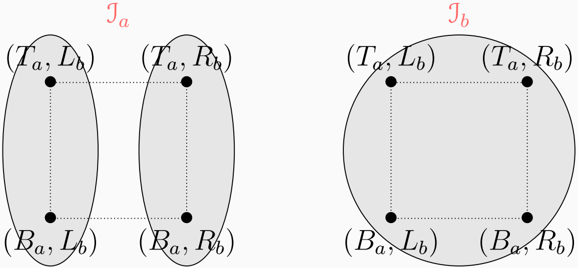

Sequential W-model

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

The closed-loop equations

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

always displays a unique solution (u_a, u_b),
whatever $\omega \in \Omega$ and W-strategies λ_a and λ_b

Solvability is a property of the information structure

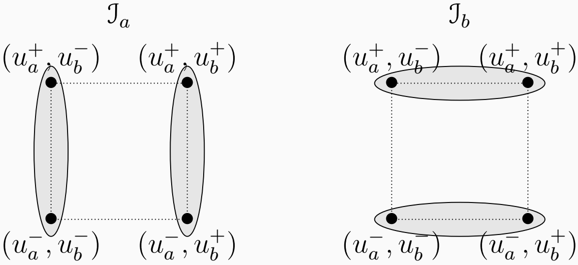

W-model with deadlock

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

The closed-loop equations

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

may display zero solutions, one solution or multiple solutions, depending on the W-strategies λ_a and λ_b

Solvability makes it possible to define a solution map from states of Nature towards configurations

Suppose that the solvability property holds true

Solution map

We define the solution map

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

that maps states of Nature towards configurations, 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 configuration space $\mathbb{H}$, making it possible to interpret $S_\lambda(\omega)$ as a configuration driven by the W-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

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

- W-strategies λ_a and λ_b have the form

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

- so that the solution map is

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

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

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

Solvable noncausal example Witsenhausen [1971]

- No Nature, $A = \{a, b, c\}$, $\mathbb{U}_a = \mathbb{U}_b = \mathbb{U}_c = \{0, 1\}$
- Set of configurations $\mathbb{H} = \{0, 1\}^3$, and information fields
  $\mathcal{I}_a = \sigma(u_b(1 - u_c))$, $\mathcal{I}_b = \sigma(u_c(1 - u_a))$, $\mathcal{I}_c = \sigma(u_a(1 - u_b))$
- The "game" can be played but… cannot be started (no first agent)

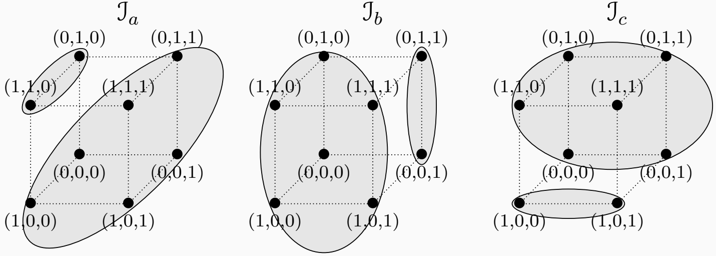

What comes next?

- Causality (as an ingredient for solvability)
- Classification of information structures

Causality

Causal configuration orderings: "Alice and Bob"

- no Nature, two agents a (Alice) and b (Bob)
- two possible actions each $\mathbb{U}_a = \{u_a^+, u_a^-\}$, $\mathbb{U}_b = \{u_b^+, u_b^-\}$
- configuration space $\mathbb{H} = \{u_a^+, u_a^-\} \times \{u_b^+, u_b^-\}$ (4 elements)
- set of total orderings (2 elements: a plays first or b plays first)
  $\Sigma^2 = \left\{ (ab) = \begin{pmatrix} \sigma:\{1,2\}\to\{a,b\} \\ \sigma(1)=a \\ \sigma(2)=b \end{pmatrix}, (ba) = \begin{pmatrix} \sigma:\{1,2\}\to\{a,b\} \\ \sigma(1)=b \\ \sigma(2)=a \end{pmatrix} \right\}$

Consider the following information structure:

- $\mathfrak{I}_b = \{\emptyset, \{u_a^+, u_a^-\}\} \otimes \{\emptyset, \{u_b^+, u_b^-\}\}$
  Bob knows nothing
- $\mathfrak{I}_a = \{\emptyset, \{u_a^+, u_a^-\}\} \otimes \{\emptyset, \{u_b^+\}, \{u_b^-\}, \{u_b^+, u_b^-\}\}$
  Alice knows what Bob does

We say that the constant configuration-ordering

- $\varphi(h) = (ab)$, for all $h \in \mathbb{H}$ (a plays first) is noncausal
- $\varphi(h) = (ba)$, for all $h \in \mathbb{H}$ (b plays first) is causal

We denote $[\![1, k]\!] = \{1, \dots, k\}$ for $k \in \mathbb{N}^*$

Partial orderings

The sets of (partial) orderings of order k are the

$$\Sigma^k = \left\{ \kappa : [\![1, k]\!] \to A \,\middle|\, \kappa \text{ is an injection} \right\}, \quad \forall k \in \mathbb{N}^*$$

The set of finite orderings is

$$\Sigma = \bigcup_{k \in \mathbb{N}^*} \Sigma^k$$

Range, cardinality, last element, first elements

For any partial ordering $\kappa \in \Sigma$, we define

the range $\|\kappa\|$ of the ordering κ as the subset of agents

$$\|\kappa\| = \big\{\kappa(1), \dots, \kappa(k)\big\} \subset A \ , \ \ \forall \kappa \in \Sigma^k$$

the cardinality $|\kappa|$ of the ordering κ as the integer

$$|\kappa| = k \in [\![1, |A|]\!] \ , \ \ \forall \kappa \in \Sigma^k$$

the last element $\kappa_\star$ of the ordering κ as the agent

$$\kappa_\star = \kappa(k) \in A \ , \ \ \forall \kappa \in \Sigma^k$$

the first elements κ_- of the ordering κ to the first $k-1$ elements

$$\kappa_- = \kappa_{|\{1,\dots,k-1\}} \in \Sigma^{k-1} \ , \ \ \forall \kappa \in \Sigma^k$$

The tree of partial orderings

There is a natural order on the set $\Sigma = \bigcup_{k \in \mathbb{N}^*} \Sigma^k$ of partial orderings

$$(\emptyset) \succeq (a) \succeq (ab) \succeq (abc)$$

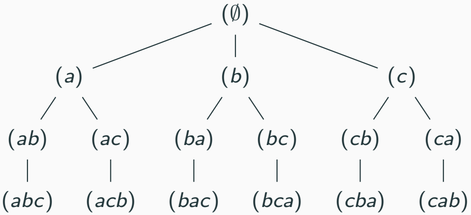

Configuration-orderings

When there is a finite or countable number $|A|$ of agents, the set of total orderings is

$$\Sigma^{|A|} = \left\{ \kappa : [\![1, |A|]\!] \to A \,\middle|\, \kappa \text{ is a bijection} \right\}$$

Configuration-ordering

A configuration-ordering is a mapping

$$\varphi : \underbrace{\mathbb{H}}_{\text{configurations}} \to \underbrace{\Sigma^{|A|}}_{\text{total orderings}}$$

The configurations $h \in \mathbb{H}$ that are compatible with a partial ordering $\kappa \in \Sigma$ belong to

$$\mathbb{H}_\kappa^\varphi = \left\{ h \in \mathbb{H} \,\middle|\, \underbrace{\varphi(h)_{|[\![1, |\kappa|]\!]}}_{\substack{\text{partial ordering of} \\ \text{the first } |\kappa| \text{ agents}}} = \kappa \right\}$$

Causality

Causal W-model

A W-model is causal if there exists (at least one) configuration-ordering $\varphi : \mathbb{H} \to \Sigma^{|A|}$ with the property that, for any $\kappa = (\kappa_-, \kappa_\star) \in \Sigma$

$$\overbrace{\underbrace{\mathbb{H}^{\varphi}_{\kappa}}_{\substack{\text{agents}\\\text{ordered by } \kappa}} \cap G}^{\substack{\text{information}\\\text{of the last agent } \kappa_\star}} \in \overbrace{\Big(\bigotimes_{c \notin \|\kappa_-\|} \{\emptyset, \mathbb{U}_c\} \otimes \bigotimes_{b \in \|\kappa_-\|} \mathcal{U}_b \otimes \mathcal{F} \Big)}^{\substack{\text{depends at most on actions}\\\text{of agents having lower order rank}}}, \quad \forall G \in \mathcal{I}_{\kappa_\star}$$

We also say that $\varphi : \mathbb{H} \to \Sigma^{|A|}$ is a causal configuration-ordering

Information comes first,
(possible) causal ordering comes second

If a W-model has no nonempty static team, it cannot be causal

A causal but nonsequential 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 action of a: thus a acts first, then b
- The reverse holds true when the state of Nature is ω^1
- A non constant configuration-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 Witsenhausen [1971]

Causality implies (recursive) solvability
with a measurable solution map

$$S_\lambda = \widetilde{S}_\lambda^{(|A|)} \circ \cdots \circ \widetilde{S}_\lambda^{(1)} \circ S_\lambda^{(0)}$$

Kuhn's extensive form of a game encapsulates causality in the tree

Solvable noncausal example Witsenhausen [1971]

- No Nature, $A = \{a, b, c\}$, $\mathbb{U}_a = \mathbb{U}_b = \mathbb{U}_c = \{0, 1\}$
- Set of configurations $\mathbb{H} = \{0, 1\}^3$, and information fields
  $\mathcal{I}_a = \sigma(u_b(1 - u_c))$, $\mathcal{I}_b = \sigma(u_c(1 - u_a))$, $\mathcal{I}_c = \sigma(u_a(1 - u_b))$
- The "game" can be played but... cannot be started (no first agent)

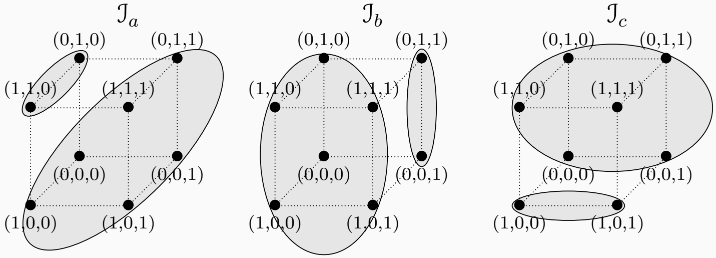

References

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

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

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*, volume 107 of *Lecture Notes in Economics and Mathematical Systems*, pages 322–335. Springer-Verlag, 1975.