Title: Von Neumannâ€•Morgenstern and Savage Theorems for Causal Decision Making

Speakers: Mauricio Soto

Series: Quantum Foundations

Date: February 29, 2024 - 11:00 AM

URL: https://pirsa.org/24020101

Abstract: Decision-making under uncertainty and causal thinking are fundamental aspects of intelligent reasoning. Decision-making has been well studied when the available information is considered at the associative (probabilistic) level. The classical Theorems of von Neumann-Morgenstern and Savage provide a formal criterion for rational choice using associative information: maximize expected utility. There is an ongoing debate around the origin of probabilities involved in such calculation. In this work, we will show how the probabilities for decision-making can be grounded in causal models by considering decision problems in which the available actions and consequences are causally connected. In this setting, actions are regarded as an intervention over a causal model. Then, we extend a previous causal decision-making result, which relies on a known causal model, to the case in which the causal mechanism that controls some environment is unknown to a rational decision-maker. In this way, action-outcome probabilities can be grounded in causal models in known and unknown cases. Finally, as an application, we extend the well-known concept of Nash Equilibrium to the case in which the players of a strategic game consider causal information.

Zoom link

Von Neumann-Morgenstern and Savage Theorems for Causal Decision Making

Mauricio Gonzalez S.

February 2024

About me: Brief CV

- BSc Mathematics - ITAM
- MSc Data Science - ITAM
- PhD Computer Science - Mexico's National Institute for Astrophysics, Optics and Electronics (INAOE)
- First Postdoc: University of Vienna

Research Interests

- Philosophy of Probability
- Foundations of Bayesian Theory
- Decision Making under Uncertainty
- Causal Inference
- Foundations of Quantum Mechanics.

Me and Physics

- How did I get involved in physics?
- Elias Okon and Daniel Sudarsky's fault...

My own philosophical position

- Because of a deeply Bayesian formation, I initially was tempted toward epistemic interpretations of QM (QBism...)
- Eventually became more of a realist.
- I still want to be able to balance realist attitudes with our position as truth-discoverers in a Kantian spirit.

Paper

- Joint work with Luis E. Sucar, David Danks, and Hugo J. Escalante.
- Paper: https://arxiv.org/abs/1907.11752

Introduction

Introduction: Causation

- An idea present since the times of Aristotle.
- Many attempts to formulate and formalize.
- Excluded from modern Statistics.
- Pearl's framework is the most used in Computer Science / AI.

Introduction: Causal Reasoning

Causal reasoning is a constant element in our lives since we constantly ask *why*:

- Why do we get sick?
- Why does a drug work?

The whole field of Causality, at least within Computer Science and AI is about taking this question seriously.

Introduction: Decision Making

On the other hand, an important aspect of acting in the world is being able to make decisions under uncertain conditions:

- Which route do we use to get to work? could there be traffic?
- Where to get some lunch? Maybe there is a long waiting line?

Von Neumann-Morgenstern

- Von Neumann and Morgenstern answered how to make choices if rational preferences are assumed and the decision maker knows the stochastic relation (i.e., probabilities of events) between actions and outcomes: **maximize expected utility**.

Von Neumann-Morgenstern

- Von Neumann and Morgenstern answered how to make choices if rational preferences are assumed and the decision maker knows the stochastic relation (i.e., probabilities of events) between actions and outcomes: **maximize expected utility**.

Choosing with objective/known probabilities

- The vNM Theorem is about choosing with objective probabilities.
- We know that thanks to De Finetti's representation theorem we can think of probabilities as frequencies, but at the same time trying to define probabilities from limiting frequencies is problematic...
- What other options do we have?

J.L. Savage

- If no such relation is known, then J.L. Savage showed that a rational decision maker must choose *as if* she is maximizing expected utility using a *subjective* probability distribution.

Página 16 de 61

Choosing with subjective probabilities

- Savage's Theorem is about choosing with subjective probabilities.
- We can turn it around and define probability as that thing used to make good decisions.

- Such theorems provide formal criteria for decision-making if rationality is assumed.

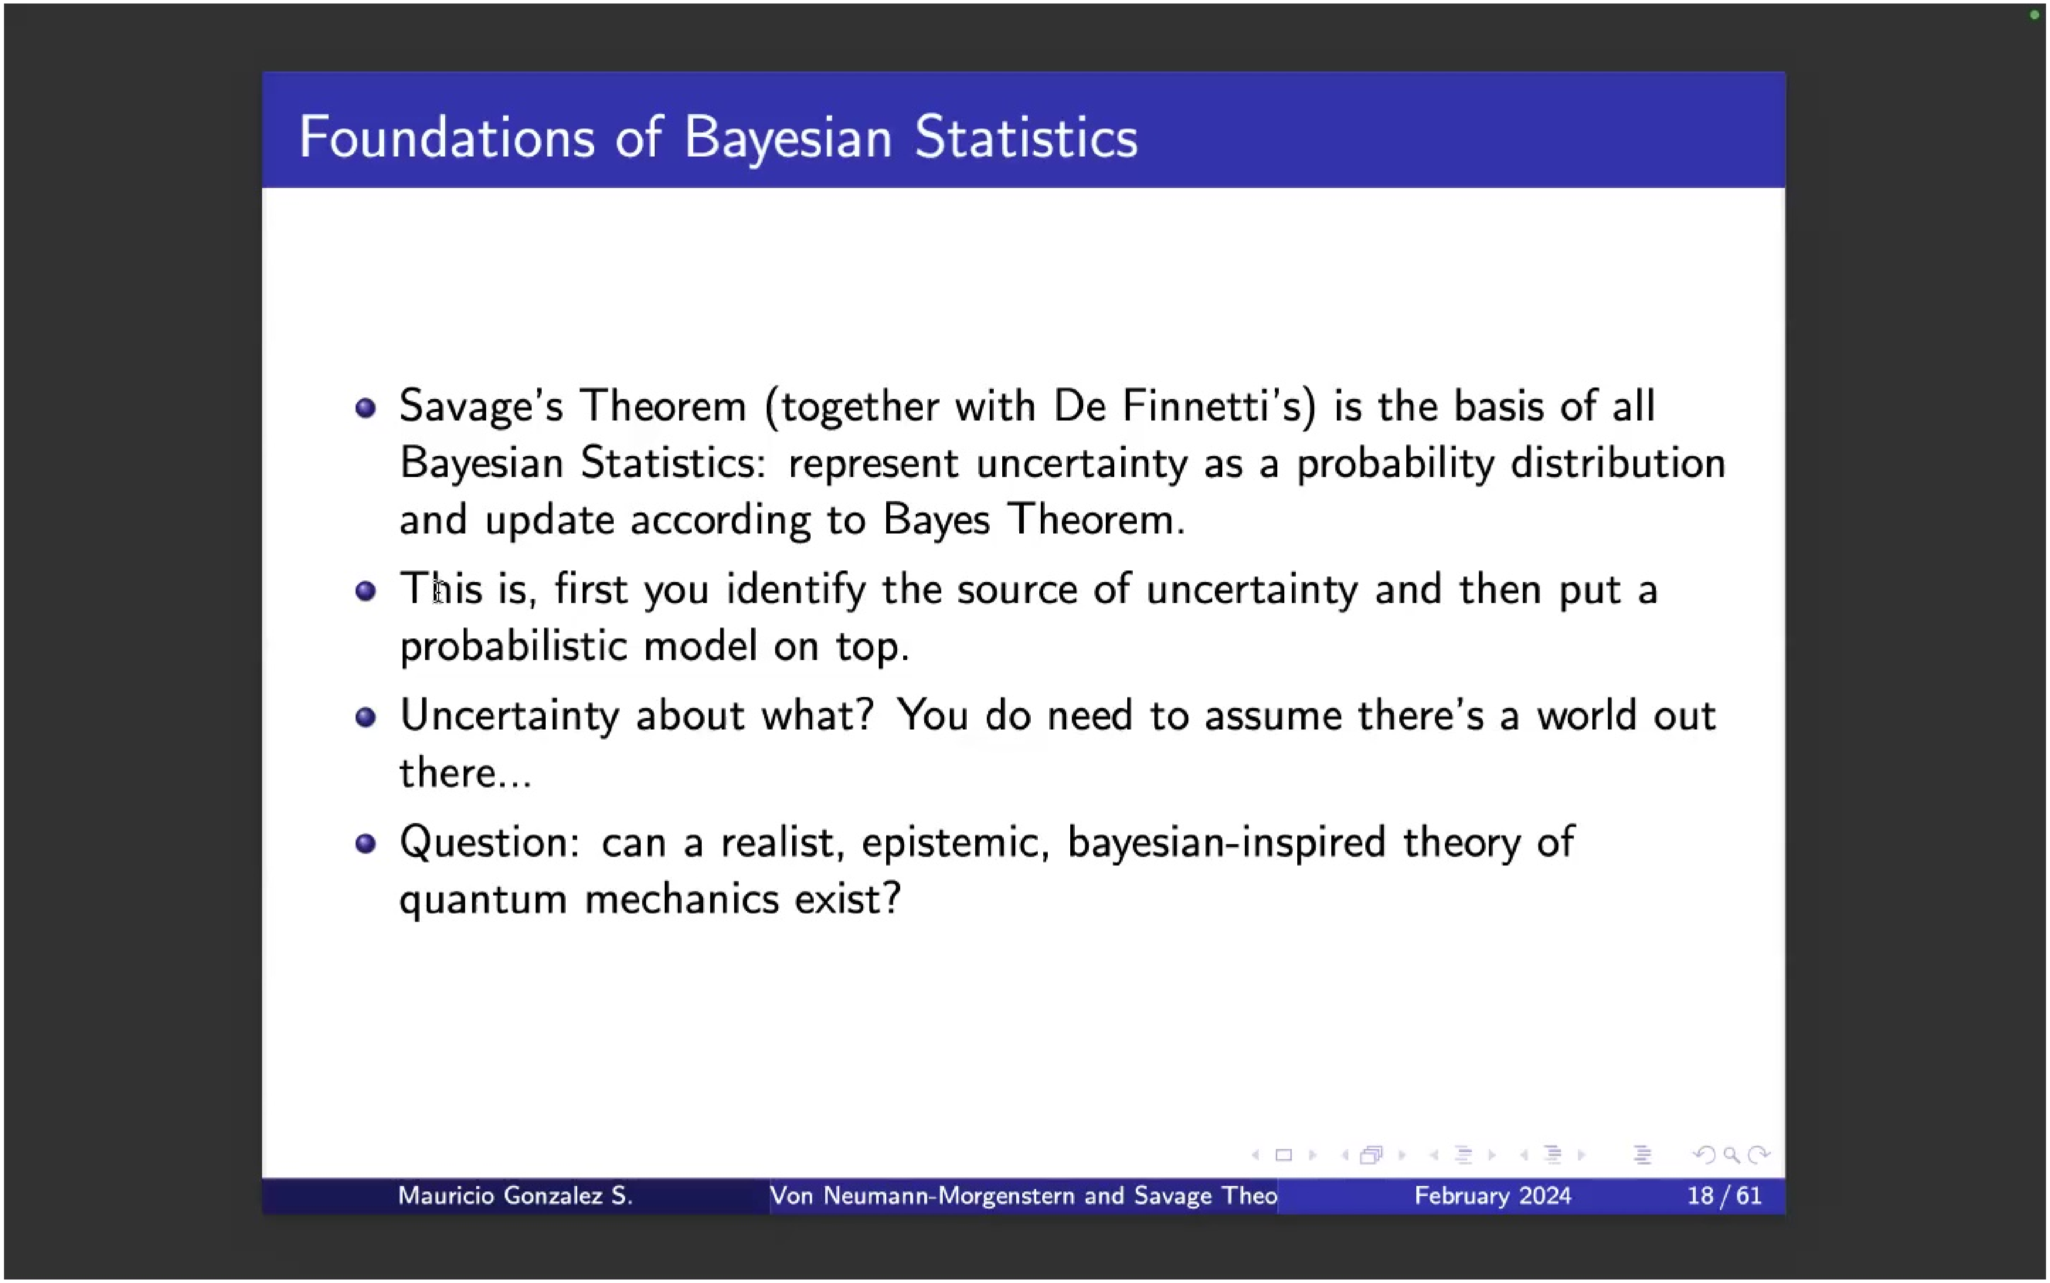
Foundations of Bayesian Statistics
Savage's Theorem (together with De Finnetti's) is the basis of all Bayesian Statistics: represent uncertainty as a probability distribution and update according to Bayes Theorem.
This is, first you identify the source of uncertainty and then put a probabilistic model on top.
Uncertainty about what? You do need to assume there's a world out there...
Question: can a realist, epistemic, bayesian-inspired theory of quantum mechanics exist?
Mauricio Gonzalez S.
Von Neumann-Morgenstern and Savage Theo
February 2024
18 / 61

Applications in Artificial Intelligence

- These criteria are the basis for many techniques used in Artificial Intelligence.
- Reinforcement Learning (RL) is a framework for learning by interaction: an agent chooses an action a, transitions to a state s, and receives reward r.
- RL algorithms attempt to learn *optimal policies*: actions prescribed by such optimal policy achieve the maximum expected utility by satisfying the Bellman Equations.

Associative data

- In general, any algorithm that relies on the von Neumann-Morgenstern or Savage Theorems, is based on *associative* relations which represent patterns found in the data in terms of correlations encoded in probability distributions.

Human causal learning and main question

- Acting in the world is conceived by human beings as *causally intervening* the world and humans can learn and use causal relations while making choices.
- A natural question is *how to formalize decision making when causal information is present?*
- Desirable to have an explicit and computationally implementable criterion for decision-making analogous to those by von Neumann-Morgenstern and Savage to provide the foundations for Causal Decision Making and Causal-based Artificial Intelligence.

Proposal

- Pearl provides an optimality criterion for decision making under causal-controlled uncertainty *when the causal mechanism which controls the environment is known* by the decision maker [7].
- We state and prove a causal decision-making theorem which considers that the causal mechanism is *unknown* to the decision maker, who is now forced to use *beliefs* about the causal relations that may hold within the variables she is considering; therefore, proposing a causal extension of Savage's Theorem.

Applications

- Our result provides the foundations for algorithms that learn best actions in uncertain causal environments, such as those found in [1, 5, 9, 2, 3].
- We will talk about some game theory.

Classical Decision-Making Theory

Decision Problems

- A Decision Problem under Uncertainty is the mathematical model of a situation where an agent must choose one of many available actions with uncertain consequences that depend on different, possibly unknown, factors.
- Such consequences are ordered in terms of the *satisfaction* that they produce in the decision maker, and such ordering is represented by a *preference relation* denoted by $\succeq$, where $a \succeq b$ is read as a being preferred to b.

Rationality

- We take the idea of rationality as a statement about the preferences of a decision-maker.
- Logically consistent.
- Taken as axioms.
- We are working within the *classical* idea of rationality.
- Other notions of rationality exist: see Prospect Theory and Behavioral Economics.

von Neumann-Morgenstern Theorem

Theorem (von Neumann-Morgenstern)

A preference relation $\succeq \subseteq L \times L$ where L is a set of lotteries with finite support over a set X satisfies the von Neumann-Morgenstern rationality axioms if and only if there exists a function $u : X \to \mathbb{R}$ such that for every $P, Q \in L$ we have that

$$P \succeq Q \text{ if and only if } \sum_{x \in X} P(x)u(x) \geq \sum_{x \in X} Q(x)u(x). \tag{1}$$

Savage's Theorem

Theorem (Savage)

In a finite, bounded Decision Problem under Uncertainty $(\mathcal{A}, \mathcal{C}, \mathcal{E}, \succeq)$, the preference relation $\succeq$ satisfies the Savage rationality axioms if and only if there exists:

- *A probability measure P, called a subjective probability, that associates with each uncertain event $E \in \mathcal{E}$ a real number $P(E)$.*
- *A utility function $u : \mathcal{C} \to \mathbb{R}$ such that it associates each consequence with a real number $u(c)$.*

Such that for a_1 and a_2 actions in $\mathcal{A}$, and any $G \neq \emptyset$

$$a_1 \succeq_G a_2$$

$$\textit{if and only if}$$

$$\sum_{j \in J(a_1)} u(c_j)P(E_j) \geq \sum_{j \in J(a_2)} u(c_j)P(E_j)$$

$$\mathbb{E}_P[u(a_1)] \geq \mathbb{E}_P[u(a_2)],$$

Limitations of the Theorems

- The mentioned theorems rely on associative data.
- This is, on correlations and patterns found in data rather than strong causal relations.
- More on this later.

Causality

- Some tens of thousands of years ago, humans began to realize that certain things cause other things.
- Tinkering with the former can change the latter

Causality

- The concept of Causality deals with regularities found in a given environment (context) that are stronger than probabilistic (or associative) relations in the sense that a causal relation allows for evaluating a change in the *consequence* given a change in the *cause*.
- This is known as an *intervention*, which is different from performing observations, and it consists of a change in the joint distribution of the variables, which is performed by forcing the value of some variable to a specific value [10, 4, 8].

Causality

- Intervention is different from observation.
- Observing a barometer falling increases the probability of a storm.
- Manually forcing barometer does not.

A purely associative world

Based only on observations:

- Patients would avoid going to the doctor to reduce the probability of being ill.
- Cities would dismiss their firefighters to reduce the number of fires.

A Definition of Causality

- Causation is a *stochastic* relation between *events* within a probability space; this is, some event (or events) *causes* another event to occur, as stated in Spirtes et al. [10].

Causal Graphical Models

- A Causal Graphical Model (CGM) consists of a set of random variables $\mathcal{X} = \{X_1, ..., X_n\}$, and a Directed Acyclic Graph (DAG) whose nodes are in correspondence with the variables in $\mathcal{X}$.
- The model is enriched with an operator named $do()$ which is defined over graphs, and whose action is described as follows:
- Given $\mathbf{X} \subseteq \mathcal{X}$ and $\mathbf{x} = \{x_{i_1}, x_{i_2}, ..., x_{i_j}\}$ an element of the set of all possible values of the variables belonging to $\mathcal{X}$, $Val(\mathcal{X})$ the action $do(\mathbf{X} = \mathbf{x})$ corresponds to assigning to each $X_j \in \mathbf{X}$ the value x_{i_j} and to delete every incoming edge into the node corresponding to each X_j in the graph $\mathcal{G}$.

Causal Decision Making

- We consider decision-making with causal information.
- There exists a causal relation between available actions and consequences in the sense that any chosen action will stochastically *cause* a consequence.

Causal Environment

Definition

A **Causal Environment** is a tuple $(\Omega, \mathcal{A}, \mathcal{G}, \mathcal{C}, \mathcal{E})$ where $(\Omega, \mathcal{A}, \mathcal{C}, \mathcal{E})$ is an uncertain environment and $\mathcal{G}$ is a CGM such that the set of uncertain events $\mathcal{E}$ correspond to the different realizations of the variables in $\mathcal{G}$ and the possible ways that the variables are related one with each other.

a von Neumann-Morgenstern Theorem for Causal Decision Making

Consider a rational decision maker who faces a causal environment in which she knows the causal model controlling the relation between her actions and outcomes. She can use the known causal model to find the probabilities of *causing* a desired outcome given she takes a certain action.

a von Neumann-Morgenstern Theorem for Causal Decision Making

Consider a rational decision maker who faces a causal environment in which she knows the causal model controlling the relation between her actions and outcomes. She can use the known causal model to find the probabilities of *causing* a desired outcome given she takes a certain action.

Pearl's decision-making result

Theorem (Causal von Neumann-Morgenstern Theorem, [7])

Let G be a Causal Graphical Model, and its associated distribution P_G. Let C be a set of consequences of interest for a decision-maker. If the decision maker faces a causal environment and if the causal graphical model G is known, then the preference relation $\succeq$ satisfies the von Neumann-Morgenstern rationality axioms if and only if:

$$a \succeq b$$

if and only if

$$\sum_{c \in C} P_G(c|do(a))u(c) \geq \sum_{c \in C} P_G(c|do(b))u(c).$$

Equivalently, the action that must be chosen is

$$a^* = \mathrm{argmax}_{a \in \mathcal{A}} \sum_{c \in C} P_G(c|do(a))u(c). \tag{2}$$

A Savage type theorem for Causal Decision Making

- We now consider **the case in which a rational decision maker does not know the causal model** which controls her environment. We must add a new axiom to the rationality axioms to keep simple our proof:
  - Choosing within a Causal Decision Problem corresponds to intervening a variable of the true causal model that controls the environment.

Main Result

Theorem (Main Result: Causal Savage Theorem)

In a finite, bounded Causal Decision Problem $(\mathcal{A}, \mathcal{G}, \mathcal{E}, \mathcal{C}, \succeq)$, where $\mathcal{G}$ is a Causal Graphical Model, we have that the preferences $\succeq$ of a decision maker are Savage-rational if and only if there exists a utility function, a probability distribution P_C over a non-empty family $\mathcal{F}$ of causal graphical models such that for each $a, b \in \mathcal{A}$:

$$a \succeq b$$
$$\textit{if and only if}$$
$$\sum_{c\in\mathcal{C}} u(c) \left(\sum_{g\in\mathcal{F}} P_g(c|do(a))P_C(g)\right)$$
$$\geq$$
$$\sum_{c\in\mathcal{C}} u(c) \left(\sum_{g\in\mathcal{F}} P_g(c|do(b))P_C(g)\right)$$

where P_g is the probability distribution associated with the causal structure g.

Interpretation

Theorem 6 says that a decision maker who faces a Causal Decision Problem is considering a probability distribution P_C over a family $\mathcal{F}$ and, within each structure, using the term $P_g(c|do(a))$ to find the probability of obtaining a certain consequence given that the intervention $do(a)$ is performed;

Application: Causal Games and Causal Nash Equilibrium

Causal Games and Causal Nash Equilibrium

- In this section we analyze an application in the domain of game theory
- We consider a *strategic game* between N rational players who are situated in a causal environment.

Games

- A game is a model of a situation in which several players must take an action and afterward they will be affected both by the outcome of their action as well as the actions of the other players [6].
- In a strategic game it is assumed that no player knows the action taken by any other players.

Bayesian Games

- We also assume that the causal mechanism, which is represented by a Causal Graphical Model $\mathcal{G}$, remains fixed and it is unknown for each player.
- In this game, players ignore the actions taken by any other player. Since the causal model that controls the environment is unknown to every player, the players also ignore the information that players will use to take their respective actions: strategic games of this type are called *Bayesian Games*.

Causal Bayesian Games

- In the games we will consider, the uncertainty of every player consists of two levels:
  - the true causal model $\mathcal{G}$
  - what an action $do(a)$ causes if a certain Causal Graphical Model ω is considered to be the true causal model.

Elements of a Bayesian Game

A Bayesian Game considers, among other things:

- For each player, a nonempty set A_i of actions.
- For each player, a finite set T_i and a function $\tau_i : \Omega \mapsto T_i$ the signal function of the player
- For each player, a probability measure p_i over Ω such that $p_i(\tau_i^{-1}(t_i)) > 0$ for all $t_i \in T_i$.

Posterior Probabilities

- We consider Ω to be a family of admissible causal models.
- $\omega \in \Omega$ being the true state of Nature fixes a causal model that controls the environment in which the players make their choices.

In classical Bayesian games, once $\omega \in \Omega$ is realized as the true state, then each player receives a signal

$$t_i = \tau_i(\omega)$$

and the posterior belief $p_i(\omega|\tau_i^{-1}(t_i))$ given by

$$p_i(\omega)/p_i(\tau_i^{-1}(t_i))$$

if $\omega \in \tau_i^{-1}(t_i)$.

In the case of causal Bayesian games:

- We must consider both the probability p_i of ω being the true state.
- As well as the probability p_i^{ω} of observing a certain consequence when doing some action a_i if ω is the true model.

Causal Nash Equilibrium

- We notice that the posterior probability itself induces a probability distribution defined over *actions* for each player once a *desired consequence* is fixed.
- This distribution, according to Theorem 6 is given by

$$p_i^{\omega}(c|do(a_i^*), a_{-i}^*)p_i(\omega|\tau_i^{-1}(t_i)).$$

Causal Nash Equilibrium

This motivates the following definition of a *Causal Nash equilibrium*:

- For each player $i \in N$ in the strategic game, we define the following probability distribution over consequences:

$$p_i^a(c) = p_i^\omega(c|do(a_i), a_{-i})p_i(\omega) \text{ for } a \in A = A_1 \times \cdots \times A_N. \qquad (3)$$

- We now define:

$$u_i^C(a) = \sum_{c \in C} u_i(c)p_i^a(c) \text{ for } a \in A = A_1 \times \cdots \times A_N. \qquad (4)$$

- Notice that u_i^C evaluates an action profile $a \in A$ in terms of the knowledge about the causal structure of each player represented by p_i, which allows each player to evaluate the probability of causing outcomes in terms of actions by using the *do* operator as well as the other actions taken by the other players, given by a_{-i} and the preferences of each player u_i.

Causal Nash Equilibrium

Using this new function, we define the equilibrium for a strategic game with causal information and Bayesian players as:

Definition

A Nash equilibrium for this *causal strategic game* is an action profile $a^* \in A$ if and only if

$$u_i^C(a^*) \geq u_i^C(a_i, a_{-i}^*) \text{ for any other } a_i \in A_i. \tag{5}$$

Causal Nash Equilibrium: interpretation

This is, an action profile is a Nash equilibrium if and only if each player uses her current knowledge about the causal structure of the environment to (causally) produce the best possible outcome given the actions taken by the other players.

Conclusion

- By extending Pearl's result to the case of unknown causal information, we have established a direct relation between choosing theorems for associative and causal information.
- Thus showing that action-outcome probabilities can be grounded in causal models.