Logic in Games

Eric Pacuit

ILLC, University of Amsterdam

staff.science.uva.nl/~epacuit

epacuit@science.uva.nl

November 29, 2006

Core Logic Lecture

If you think that your paper is vacuous,
Use the first-order functional calculus.
It then becomes logic,
And, as if by magic,
The obvious is hailed as miraculous.

(Moshe Vardi)

What Logic in Which Games?

"Game theory is a bag of analytical tools designed to help us understand the phenomena that we observe when decision-makers interact."

Osborne and Rubinstein. *Introduction to Game Theory.* MIT Press .

What Logic in Which Games?

Game theory is full of deep puzzles, and there is often disagreement about proposed solutions to them. The puzzlement and disagreement are neither empirical nor mathematical but, rather, concern the meanings of fundamental concepts ('solution', 'rational', 'complete information') and the soundness of certain arguments...Logic appears to be an appropriate tool for game theory both because these conceptual obscurities involve notions such as reasoning, knowledge and counter-factuality which are part of the stock-in-trade of logic, and because it is a prime function of logic to establish the validity or invalidity of disputed arguments.

M.O.L. Bacharach. *Logic and the Epistemic Foundations of Game Theory.* .

(Modal) Logic in Games

M. Pauly and W. van der Hoek. *Modal Logic form Games and Information.* Handbook of Modal Logic (2006).

G. Bonanno. *Modal logi and game theory: Two alternative approaches.* Risk Decision and Policy **7** (2002).

J. van Benthem. *Extensive games as process models.* Journal of Logic, Language and Information **11** (2002).

J. Halpern. *A computer scientist looks at game theory.* Games and Economic Behavior **45:1** (2003).

R. Parikh. *Social Software.* Synthese **132: 3** (2002).

Logic in Games: Relevant Conferences

LOFT: Conference on Logic and the Foundations of Game and Decision Theory (Amsterdam: `www.illc.uva.nl/LOFT2008/`)

TARK: Theoretical Aspects of Rationality and Knowledge (Brussels 2007: `www.info.fundp.ac.be/ pys/TARK07/`)

GLoRiClass Seminar: `www.illc.uva.nl/GLoRiClass/`

New perspectives on Games and Interaction, Feb. 5 - 7, 2007 (`www.illc.uva.nl/KNAW-AC/`)

What I Want to Talk About

- Game Logics
- Logics for social interactive situations
- When are two games the *same*?
- Epistemic program in game theory
- Aggregating individual judgments
- Axiomatization results in Social Choice
- (Formally) Verifying that a social procedure is *correct*
- Social Software
  =(Social Choice+Game Theory+Computer Science)/Logic
- Develop ("well-behaved") logical languages that can express game theoretic concepts, such as the *Nash equilibrium*

What I Will Talk About

- Game Logics
- Logics for social interactive situations
- When are two games the *same*?
- Epistemic program in game theory
- Aggregating individual judgments
- Axiomatization results in Social Choice
- (Formally) Verifying that a social procedure is *correct*
- Social Software
  =(Social Choice+Game Theory+Computer Science)/Logic
- Develop ("well-behaved") logical languages that can express game theoretic concepts, such as the *Nash equilibrium*

What I Will *Actually* Talk About

- A paradox surrounding the epistemic foundations of solution concepts
- An “Axiomatization” result in Social Choice Theory
- Some impossibility results
- A “logical approach” to backwards induction

Goal: Illustrate where logic naturally shows up in the social sciences and point to some relevant literature.

Epistemic Program in Game Theory

Fundamental Problem: What does it mean to say that the players in a game are **rational**, each **thinks** each other is rational, each thinks each other thinks the others are rational, etc.?

Epistemic Program in Game Theory

Fundamental Problem: What does it mean to say that the players in a game are **rational**, each **thinks** each other is rational, each thinks each other thinks the others are rational, etc.?

Epistemic Program in Game Theory: An explicit description of the players' beliefs is part of the descripiton of a game.

Identify for any game the strategies that are chosen by rational and intelligent players who know the strucutre of the game, the preference of the other players and recognize each others rationality and beliefs.

Literature

See, for example,

R. Aumann. *Interactive Epistemology I & II.* International Journal of Game Theory (1999).

P. Battigalli and G. Bonanno. *Recent results on belief, knowledge and the epistemic foundations of game theory.* Research in Economics (1999).

B. de Bruin. *Explaining Games.* Ph.D. Thesis, 2004.

R. Stalnaker. *Belief Revision in Games: Forward and Backward Induction.* Mathematical Social Sciences (1998).

Describing Beliefs

Fix a set of possible states (complete descriptions of a situation). Two main approaches to describe beleifs:

- Set-theortical (Kripke Structures, Aumann Structures): For each state and each agent i, specify a set of states that i considers possible.

- Probabilistic (Bayesian Models, Harsanyi Type Spaces): For each state, define a (subjective) probability function over the set of states for each agent.

A Question

Which model is the “correct” one to work with?

A Question

Which model is the “correct” one to work with?

It turns out that finding the connection between rationality, what agents think about the situation and what actually happens depends on the existence of a “rich enough” space of types, i.e., a universal type space.

A Question

Which model is the "correct" one to work with?

It turns out that finding the connection between rationality, what agents think about the situation and what actually happens depends on the existence of a "rich enough" space of types, i.e., a universal type space.

> *It is not enough [...] that Ann should consider each of Bob's strategies possible. Rather, she considers possible both every strategy that Bob might play and every type that Bob might be. (Likewise, Bob considers possible both every strategy that Ann might play and every type that Ann might be.)*

Brandenburger, Friedenburg and Keisler. *Admissibility in Games.* 2004.

A Paradox

Ann believes that Bob assumes* that
Ann believes that Bob's assumption is wrong.

Does Ann believe that Bob's assumption is wrong?

* An **assumption** (or strongest belief) is a belief that implies all other beliefs.

A. Brandenburger and H. J. Keisler. *An Impossibility Theorem on Beliefs in Games.* forthcoming in *Studia Logica.*

A Paradox

Ann believes that Bob assumes* that
Ann believes that Bob's assumption is wrong.

Does Ann believe that Bob's assumption is wrong? Yes.

A Paradox

Ann believes that Bob assumes* that
Ann believes that Bob's assumption is wrong.

Does Ann believe that Bob's assumption is wrong? Yes.

Then according to Ann, Bob's assumption is wrong.

A Paradox

Ann believes that Bob assumes* that
Ann believes that Bob's assumption is wrong.

Does Ann believe that Bob's assumption is wrong? Yes.

Then according to Ann, Bob's assumption is wrong.

A Paradox

Ann believes that Bob assumes* that
Ann believes that Bob's assumption is wrong.

Does Ann believe that Bob's assumption is wrong? Yes.

Then according to Ann, Bob's assumption is wrong.

So, according to Ann, Bob's assumption is correct — i.e., *Bob's assumption is not wrong.*

A Paradox

Ann believes that Bob assumes* that
Ann believes that Bob's assumption is wrong.

Does Ann believe that Bob's assumption is wrong? Yes.

Then according to Ann, Bob's assumption is wrong.

So, according to Ann, Bob's assumption is correct — i.e., *Bob's assumption is not wrong.*

So, the answer must be no.

A Paradox

Ann believes that Bob assumes* that
Ann believes that Bob's assumption is wrong.

Does Ann believe that Bob's assumption is wrong? No.

A Paradox

Ann believes that Bob assumes* that
Ann believes that Bob's assumption is wrong.

Does Ann believe that Bob's assumption is wrong? No.

Then according to Ann, Bob's assumption is correct.

A Paradox

Ann believes that Bob assumes* that
Ann believes that Bob's assumption is wrong.

Does Ann believe that Bob's assumption is wrong? No.

Then according to Ann, Bob's assumption is correct.

A Paradox

Ann believes that Bob assumes* that
Ann believes that Bob's assumption is wrong.

Does Ann believe that Bob's assumption is wrong? No.

Then according to Ann, Bob's assumption is correct.

That is, it is correct that Ann believes that Bob's assumption is wrong.

A Paradox

Ann believes that Bob assumes* that
Ann believes that Bob's assumption is wrong.

Does Ann believe that Bob's assumption is wrong? No.

Then according to Ann, Bob's assumption is correct.

That is, it is correct that Ann believes that Bob's assumption is wrong.

So, the answer must be yes.

Main Result

Belief Model: a set of states for each player, and a relation for each player that specifies when a state of one player considers a state of the other player to be possible.

Main Result

Belief Model: a set of states for each player, and a relation for each player that specifies when a state of one player considers a state of the other player to be possible.

Language: the language used by the players to formulate their beliefs

Main Result

Belief Model: a set of states for each player, and a relation for each player that specifies when a state of one player considers a state of the other player to be possible.

Language: the language used by the players to formulate their beliefs

Complete: A belief model is complete for a language if every statement in a player's langauge which is possible (i.e. true for some states) can be assumed by the player.

Main Result

Belief Model: a set of states for each player, and a relation for each player that specifies when a state of one player considers a state of the other player to be possible.

Language: the language used by the players to formulate their beliefs

Complete: A belief model is complete for a language if every statement in a player's langauge which is possible (i.e. true for some states) can be assumed by the player.

Theorem (Brandenburger and Keisler) No belief model can be complete for a language that contains first-order logic.

Open Question

Can we find a logic $\mathcal{L}$ such that

1. Complete belief models for $\mathcal{L}$ exist for each game;

2. notions such as rationality, belief in rationality, etc. are expressible in $\mathcal{L}$; and

3. the ingredients in 1 and 2 can be combined to yield various well-known game-theoretic solution concepts.

Aggregating Preferences: Some Notation

- Suppose that there are n individuals and two alternatives x and y

- Let xP_iy denote that i prefers x to y and xI_iy denote that i is indifferent between x and y

Aggregating Preferences: Some Notation

- For each i there is a variable $D_i \in \{-1, 0, 1\}$ where

$$D = \begin{cases} -1 & \text{if } \; yP_ix \\ 0 & \text{if } \; xI_iy \\ 1 & \text{if } \; xP_iy \end{cases}$$

- $f : \{-1, 0, 1\}^n \rightarrow \{-1, 0, 1\}$ is the *group decision function*

Simple Majority Procedure

For $k \in \{-1, 0, 1\}$, let

$$N_k(D_1, \ldots, D_n) = |\{i \mid D_i = k\}|$$

Let $\vec{D} = \langle D_1, \ldots, D_n \rangle$

f is a **simple majority decision method** iff

$$f(\vec{D}) = \begin{cases} -1 & \text{if } N_1(\vec{D}) - N_{-1}(\vec{D}) < 0 \\ 0 & \text{if } N_1(\vec{D}) - N_{-1}(\vec{D}) = 0 \\ 1 & \text{if } N_1(\vec{D}) - N_{-1}(\vec{D}) > 0 \end{cases}$$

Properties of group decision funcitons

A group decision function f is

- **Decisive** if it is a total function

- **Symmetric** if $f(D_1, \ldots, D_n) = f(D_{j(1)}, \ldots, D_{j(n)})$ for all permutations j. I.e., f is symmetric in all of its arguments.

- **Neutral** if $f(-D_1, \ldots, -D_n) = -f(D_1, \ldots, D_n)$

- **Positively Responsive** if $D = f(D_1, \ldots, D_n) = 1/2$ or 1, and $D'_i = D_i$ for all $i \neq i_0$, and $D'_{i_0} > D_{i_0}$, then $D' = f(D'_1, \ldots, D'_n) = 1$

May's Theorem

Theorem (May, 1952) A group decision function is the method of simple majority decision if and only if it is always decisive, symmetric, neutral and positively responsive

May's Theorem

Theorem (May, 1952) A group decision function is the method of simple majority decision if and only if it is always decisive, symmetric, neutral and positively responsive

Formal Minimalism

M. Pauly. *On the Role of Language in Social Choice Theory*. Available at the author's website (2005).

Generalizing May's Theorem

In May's Theorem, the agents are making a single binary choice between two alternatives. What about more general situations?

Generalizing May's Theorem

In May's Theorem, the agents are making a single binary choice between two alternatives. What about more general situations?

- Agents choose between between more than two alternatives.
- There are multiple interconnected propositions on which simultaneous decisions are to be made.

Condorcet Paradox

Suppose that there are three agents choosing between three alternatives.

$$\begin{array}{ll} P_1 & a > b > c \\ P_2 & b > c > a \\ P_3 & c > a > b \end{array}$$

Pairwise majority voting produces a non-transitive group preference.

$$\begin{array}{ll} P_1 & a > b > c \\ P_2 & b > c > a \\ P_3 & c > a > b \end{array}$$

- $a > b$?
- $b > c$?
- $a > c$?

$$\begin{array}{ll} P_1 & a > b > c \\ P_2 & b > c > a \\ P_3 & c > a > b \end{array}$$

- $a > b$?
- $b > c$?
- $a > c$?

$$
\begin{array}{ll}
P_1 & a > b > c \\
P_2 & b > c > a \\
P_3 & c > a > b
\end{array}
$$

- $a > b$? Yes
- $b > c$?
- $a > c$?

P_1	$a > b > c$
P_2	$b > c > a$
P_3	$c > a > b$

- $a > b$? Yes
- $b > c$? Yes
- $a > c$?

P_1	$a > b > c$
P_2	$b > c > a$
P_3	$c > a > b$

- $a > b$? Yes
- $b > c$? Yes
- $a > c$? No

Arrow's Theorem: Some Notation

Let $\mathcal{R}$ be the set of all reflexive, transitive and connected relations on a set of candidates X.

A **social welfare function** F is a function $F : \mathcal{R}^n \to \mathcal{R}$

Suppose that $R = F(R_1, \ldots, R_n)$

Arrow's Theorem: Conditions

- **Universal Domain**: F is a total function
- **Weak Pareto Principle**: For any two candidates x, y if xR_iy for each agent i then $xF(\vec{R})y$
- **Independence of Irrelevant Alternatives**: Suppose that $\vec{R}$ and $\vec{R}^*$ are two preference profiles and x and y are two candidates such that for all individuals i, if xR_iy iff xR_i^*y then $xF(\vec{R})y$ iff $xF(\vec{R}^*)y$.
- **Non-Dictatorship**: There does not exist an individual i such that for all profiles $\vec{R} \in \mathcal{R}^n$, if xR_iy then $xF(\vec{R})y$.

Arrow's Theorem

Theorem (Arrow 1951/1963) There exists no social welfare function which satisfies Universal Domain, Weak Pareto Principle, Independence of Irrelevant Alternatives and Non-Dictatorship.

Arrow's Theorem

Theorem (Arrow 1951/1963) There exists no social welfare function which satisfies Universal Domain, Weak Pareto Principle, Independence of Irrelevant Alternatives and Non-Dictatorship.

Formalizing Arrow's Theorem

T. Agotnes, W. van der Hoek and M. Wooldridge. *Towards a Logic of Social Welfare*. Proceedings of LOFT, COMSOC (2006).

The Doctrinal Paradox

P: “UvA teachers get a 10% raise"

Q: “The quality of education for all students will increase"

$P \to Q$: “If UvA teachers get a 10% raise, then the quality of education for all students will increase"

	P	$P \to Q$	Q
Individual 1	True	True	True
Individual 2	True	False	False
Individual 3	False	True	False
Majority	True	True	False

A Second Paradox (Kornhauser and Sager 1993)

P: a valid contract was in place

Q: the defendant's behaviour was such as to breach a contract of that kind

R: the court is required to find the defendant liable.

	P	Q	$(P \wedge Q) \leftrightarrow R$	R
1	yes	yes	yes	yes
2	yes	no	yes	no
3	no	yes	yes	no

Should we accept R?

	P	Q	$(P \wedge Q) \leftrightarrow R$	R
1	yes	yes	yes	yes
2	yes	no	yes	no
3	no	yes	yes	no

Should we accept R? No, a simple majority votes no.

	P	Q	$(P \wedge Q) \leftrightarrow R$	R
1	yes	yes	yes	yes
2	yes	no	yes	no
3	no	yes	yes	no

Should we accept R? Yes, a majority votes yes for P and Q and $(P \wedge Q) \leftrightarrow R$ is a legal doctrine.

	P	Q	$(P \wedge Q) \leftrightarrow R$	R
1	yes	yes	yes	yes
2	yes	no	yes	no
3	no	yes	yes	no

List and Pettit Impossibility Result

Suppose there are n agetns and let $\mathcal{L}$ be a propositional language.

Personal judgement sets: a *consistent*, *complete* and *deductively closed* set of formulas — a maximally consistent set.

A **collective judgement aggregation function:** Let $\mathbb{M} = \{\Gamma \mid \Gamma \text{ is a maximally consistent set}\}$ then a collective aggregation function is defined as follows:

$$F : \mathbb{M}^n \to \mathbb{M}$$

Some Conditions

Universal Domain F is a total function

Anonymity For all $\vec{\Gamma} \in \mathbb{M}^n$, $F(\Gamma_1, \ldots, \Gamma_n) = F(\Gamma_{\pi(1)}, \ldots, \Gamma_{\pi(n)})$ for all permutations π

Systematicity There exists a function $f : \{0,1\}^n \rightarrow \{0,1\}$ such that for any $\vec{\Gamma} \in \mathbb{M}^n$,
$F(\Gamma_1, \ldots, \Gamma_n) = \{\phi \in X \mid f(\delta_1(\phi), \ldots, \delta_n(\phi)) = 1\}$, where, for each agent i and each $\phi \in X$, $\delta_i(\phi) = 1$ if $\phi \in \Gamma_i$ and $\delta_i(\phi) = 0$ if $\phi \notin \Gamma_i$

Theorem (List and Pettit, 2001) There exists no judgement aggregation function generating complete, consistent and deductively closed collective sets of judgements which satisfies Universal Domain, Anonymity and Systematicity.

Brief Survey of the Literature

- See `personal.lse.ac.uk/LIST/doctrinalparadox.htm` for a detailed overview of the current state of affairs. Some highlights:
- Other impossibility results: Pauly and van Hees (2003), van Hees (2004), Gärdenfors (2004), and others
- List and Pettit (2005) compare their impossibility result with Arrow's Theorem
- For a general approach see Daniëls and Pacuit (2006).

Backward Induction

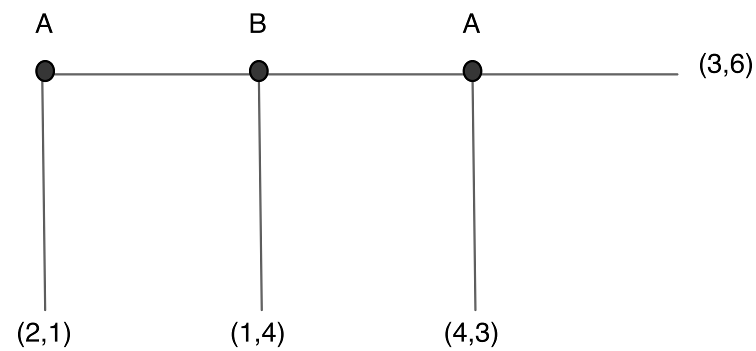

Backward Induction

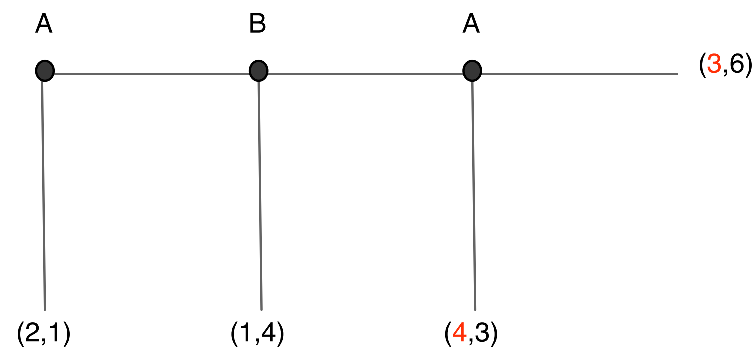

Backward Induction

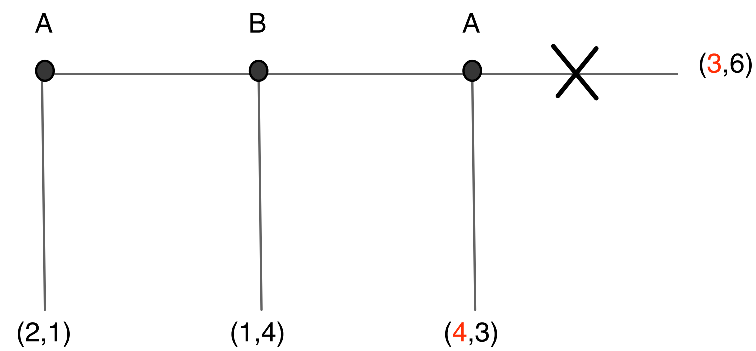

Backward Induction

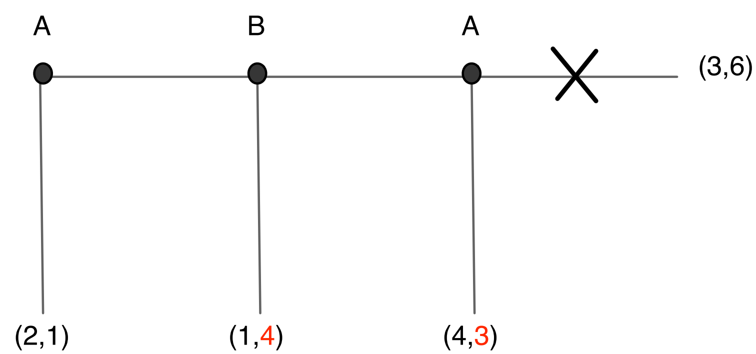

Backward Induction

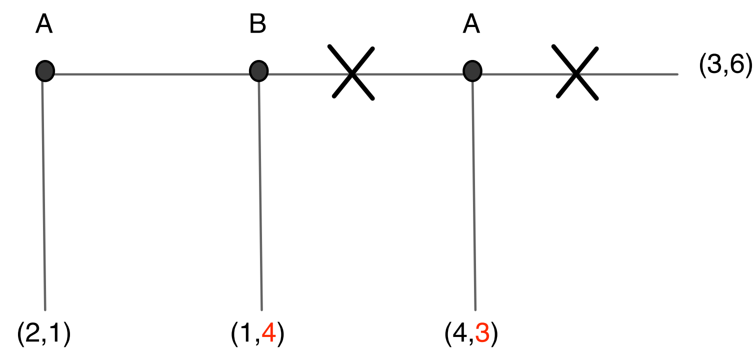

Backward Induction

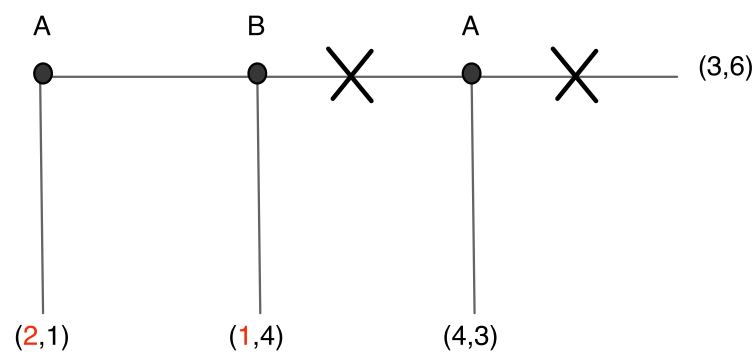

Backward Induction

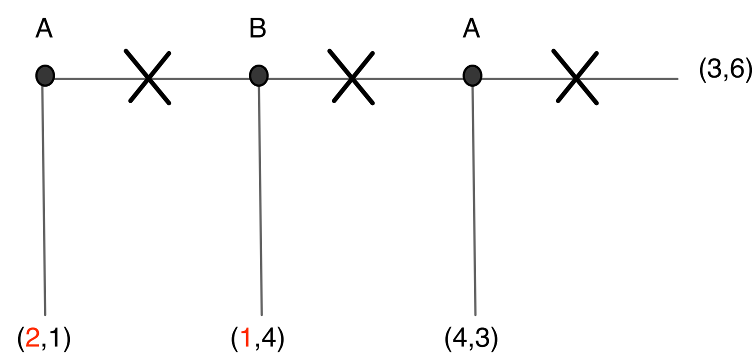

Backward Induction

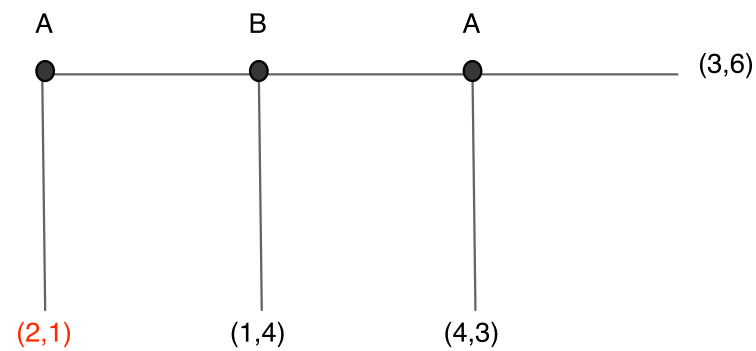

Backward Induction

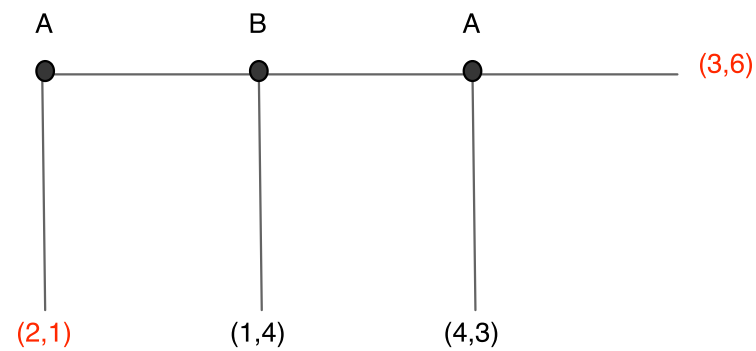

A Logical Characterization of Backwards Induction

Models: Extensive games (labeled trees with preference relations over the end nodes)

Goal: find a language and a formula from that language that "characterizes" the backward induction relation.

A Logical Characterization of Backwards Induction

What do we want to express?

- "after action a, ϕ is true": $\langle a \rangle \phi$

A Logical Characterization of Backwards Induction

What do we want to express?

- "after action a, ϕ is true": $\langle a \rangle \phi$
- "agent i's turn to move": $turn_i$

A Logical Characterization of Backwards Induction

What do we want to express?

- "after action a, ϕ is true": $\langle a \rangle \phi$
- "agent i's turn to move": $turn_i$
- "after some move ϕ is true": $\langle move \rangle \phi$

A Logical Characterization of Backwards Induction

What do we want to express?

- "after action a, ϕ is true": $\langle a \rangle \phi$
- "agent i's turn to move": $turn_i$
- "after some move ϕ is true": $\langle move \rangle \phi$
- "ϕ is true after the agent chooses in its best interest": $\langle bi \rangle \phi$

A Logical Characterization of Backwards Induction

What do we want to express?

- "after action a, ϕ is true": $\langle a \rangle \phi$
- "agent i's turn to move": $turn_i$
- "after some move ϕ is true": $\langle move \rangle \phi$
- "ϕ is true after the agent chooses in its best interest": $\langle bi \rangle \phi$
- "ϕ is true in a *preferred node*": $\Diamond_i \phi$

A Logical Characterization of Backwards Induction

What do we want to express?

- "after action a, ϕ is true": $\langle a \rangle \phi$
- "agent i's turn to move": $turn_i$
- "after some move ϕ is true": $\langle move \rangle \phi$
- "ϕ is true after the agent chooses in its best interest": $\langle bi \rangle \phi$
- "ϕ is true in a *preferred node*": $\Diamond_i \phi$
- "ϕ is true after agents *repeatedly* choose in their best interests" $\langle bi^* \rangle \phi$

A Logical Characterization of Backwards Induction

Proposition The relation bi corresponding to a unique outcome of a Backward Induction computation is the only binary relation on a game model satisfying the following principles for all propositions ϕ:

1. $\langle move \rangle \top \to (\langle bi \rangle \neg\phi \to \neg\langle bi \rangle \phi)$
2. For all players i,
   $(turn_i \wedge \langle bi^* \rangle (end \wedge \phi)) \to [move]\langle bi^* \rangle (end \wedge \Diamond_i \phi)$

J. van Benthem, S. van Otterloo and O. Roy. *Preference logic, conditionals and solution concepts in games*. ILLC Prepublications 2005.

Conclusion

What can Logic do for Game Theory?

See `staff.science.uva.nl/~epacuit/caputLLI.html` for more information.

Thank you.

Verifying Social Procedures

M. Pauly and M. Wooldridge. *Logic for Mechanism Design — A Manifesto*. Available at authors website (2005).

Computational vs. Behavioral Structures

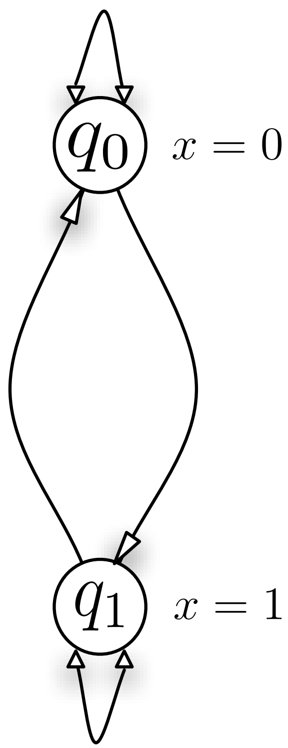

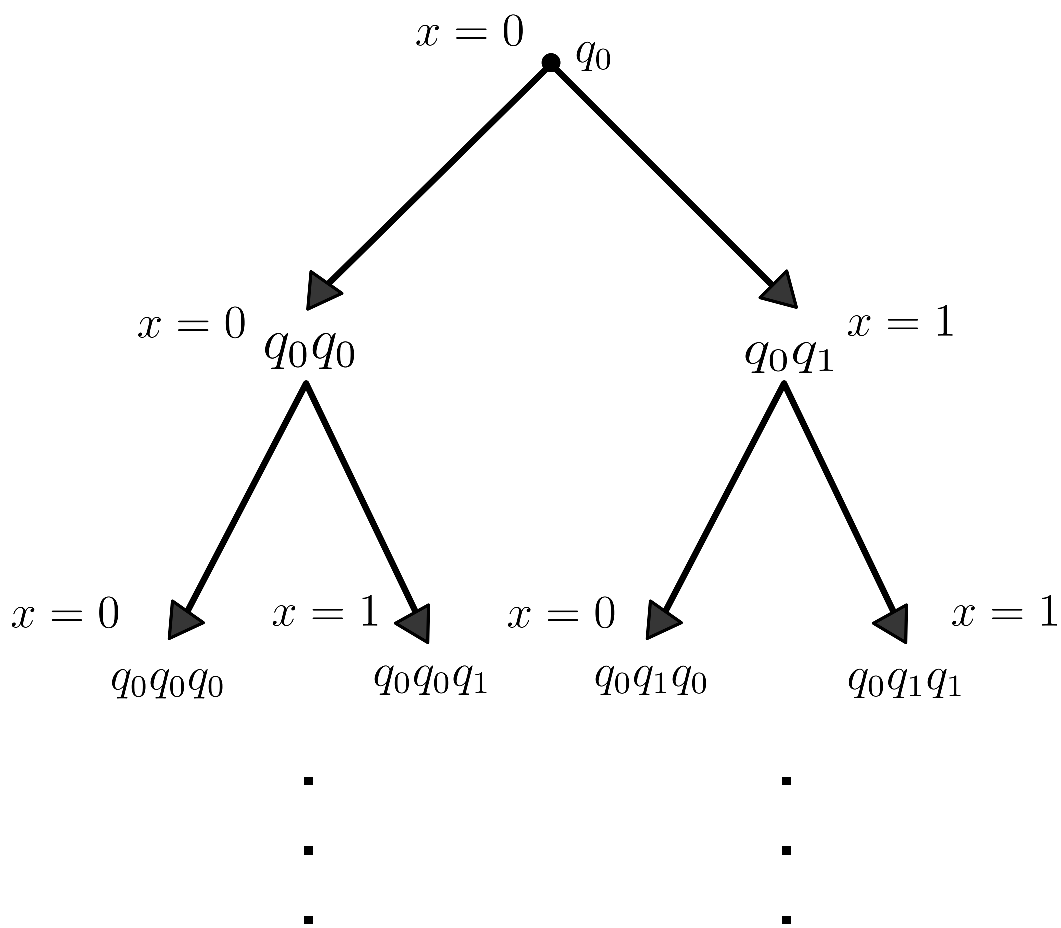

$$\neg\forall\Diamond P_{x=1}$$

Alternating Transition Systems

The previous model assumes there is *one* agent that "controls" the transition system.

Alternating Transition Systems

The previous model assumes there is *one* agent that "controls" the transition system.

What if there is more than one agent?

Alternating Transition Systems

The previous model assumes there is *one* agent that “controls” the transition system.

What if there is more than one agent?

Example: Suppose that there are two agents: a server (s) and a client (c). The client asks to set the value of x and the server can either grant or deny the request. Assume the agents make simultaneous moves.

Alternating Transition Systems

The previous model assumes there is *one* agent that "controls" the transition system.

What if there is more than one agent?

Example: Suppose that there are two agents: a server (s) and a client (c). The client asks to set the value of x and the server can either grant or deny the request. Assume the agents make simultaneous moves.

	$deny$	$grant$
$set0$		
$set1$		

Alternating Transition Systems

The previous model assumes there is *one* agent that "controls" the transition system.

What if there is more than one agent?

Example: Suppose that there are two agents: a server (s) and a client (c). The client asks to set the value of x and the server can either grant or deny the request. Assume the agents make simultaneous moves.

	$deny$	$grant$
$set0$		$q_0 \Rightarrow q_0$, $q_1 \Rightarrow q_0$
$set1$		$q_0 \Rightarrow q_1$, $q_1 \Rightarrow q_1$

Alternating Transition Systems

The previous model assumes there is *one* agent that "controls" the transition system.

What if there is more than one agent?

Example: Suppose that there are two agents: a server (s) and a client (c). The client asks to set the value of x and the server can either grant or deny the request. Assume the agents make simultaneous moves.

	$deny$	$grant$
$set0$	$q \Rightarrow q$	$q_0 \Rightarrow q_0,\ q_1 \Rightarrow q_0$
$set1$	$q \Rightarrow q$	$q_0 \Rightarrow q_1,\ q_1 \Rightarrow q_1$

Multi-agent Transition Systems

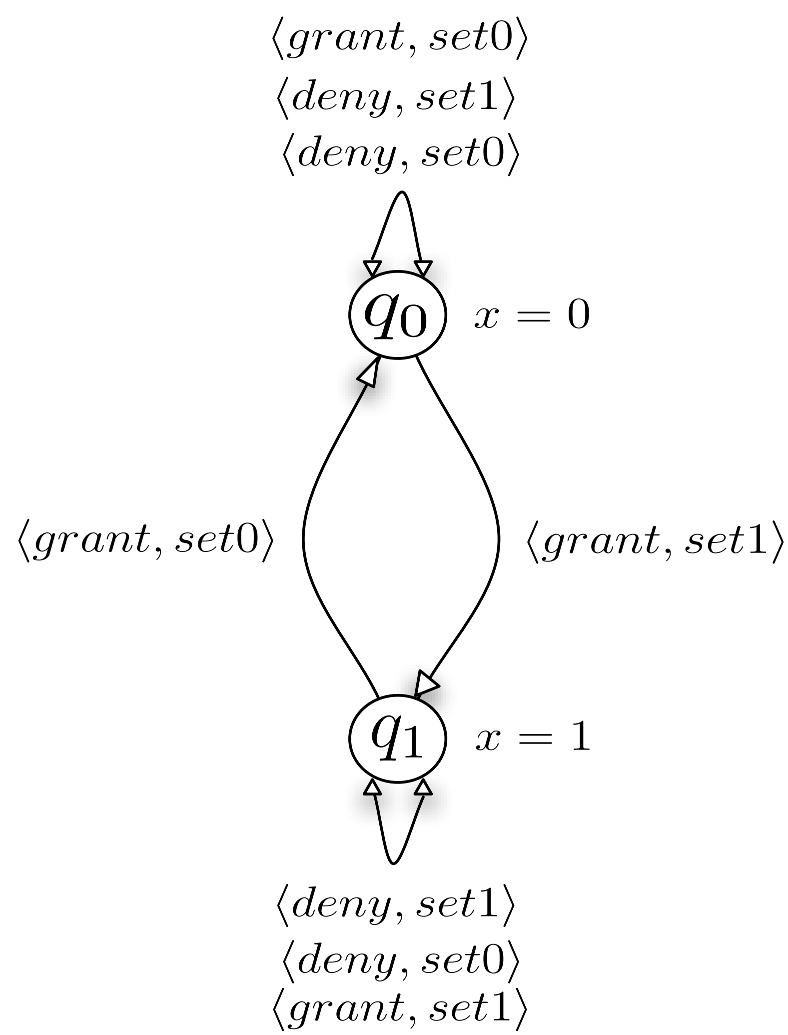

Multi-agent Transition Systems

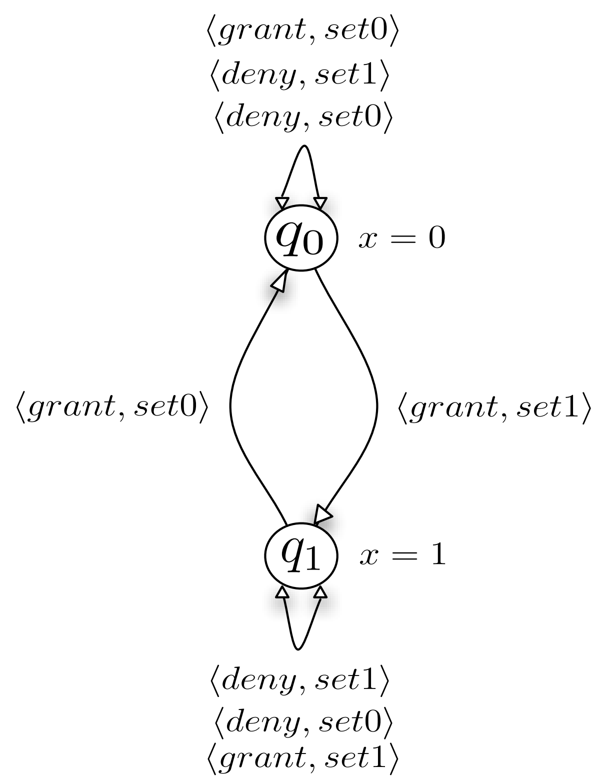

$$(P_{x=0} \to [s]P_{x=0}) \wedge (P_{x=1} \to [s]P_{x=1})$$

Multi-agent Transition Systems

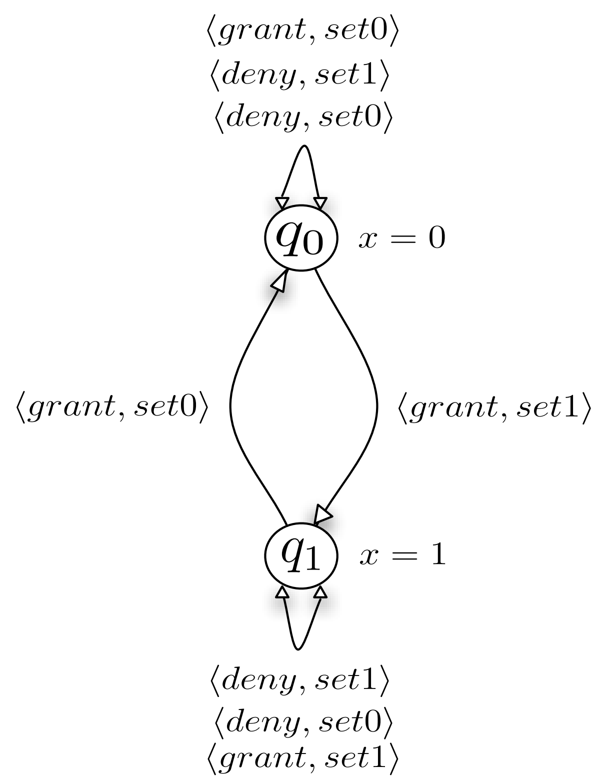

$$P_{x=0} \longrightarrow \neg[s]P_{x=1}$$

Multi-agent Transition Systems

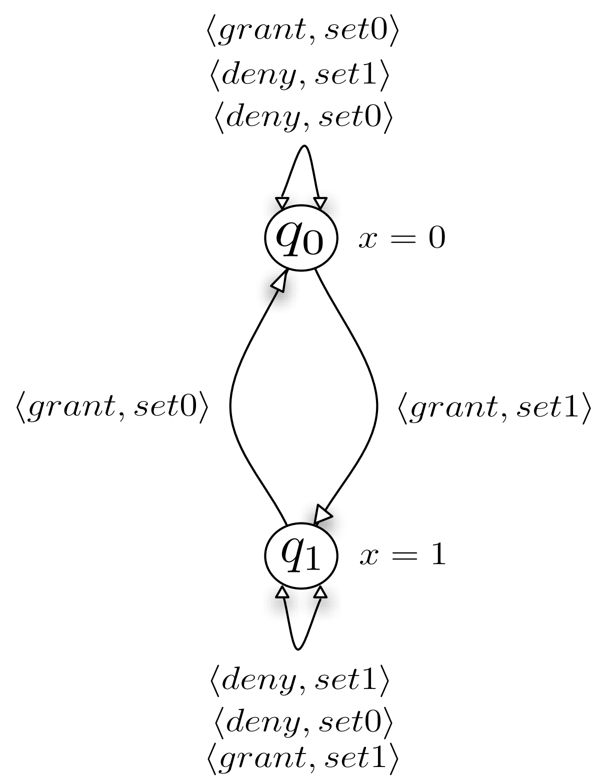

$$P_{x=0} \rightarrow [s, c]P_{x=1}$$

From Temporal Logic to Strategy Logic

From Temporal Logic to Strategy Logic

- *Linear Time Temporal Logic*: Reasoning about computation paths:

  $\Diamond\phi$: ϕ is true some time in *the* future.

A. Pnuelli. *A Temporal Logic of Programs.* in *Proc. 18th IEEE Symposium on Foundations of Computer Science* (1977).

From Temporal Logic to Strategy Logic

- *Linear Time Temporal Logic*: Reasoning about computation paths:

$$\Diamond\phi\text{: } \phi \text{ is true some time in } \textit{the} \text{ future.}$$

A. Pnuelli. *A Temporal Logic of Programs.* in *Proc. 18th IEEE Symposium on Foundations of Computer Science* (1977).

- *Branching Time Temporal Logic*: Allows quantification over paths:

$$\exists\Diamond\phi\text{: there is a path in which } \phi \text{ is eventually true.}$$

E. M. Clarke and E. A. Emerson. *Design and Synthesis of Synchronization Skeletons using Branching-time Temproal-logic Specifications.* In *Proceedings Workshop on Logic of Programs*, LNCS (1981).

From Temporal Logic to Strategy Logic

- *Alternating-time Temporal Logic*: Reasoning about (local and global) group power:

  $\langle\langle A\rangle\rangle\Box\phi$: The coalition A has a joint strategy to ensure that ϕ will remain true.

R. Alur, T. Henzinger and O. Kupferman. *Alternating-time Temproal Logic*. *Jouranl of the ACM* (2002).

From Temporal Logic to Strategy Logic

- *Alternating-time Temporal Logic*: Reasoning about (local and global) group power:

  $\langle\langle A\rangle\rangle\Box\phi$: The coalition A has a joint strategy to ensure that ϕ will remain true.

R. Alur, T. Henzinger and O. Kupferman. *Alternating-time Temproal Logic. Jouranl of the ACM* (2002).

- *Coalitional Logic*: Reasoning about (local) group power (fragment of **ATL**).

  $[C]\phi$ (equivalently $\langle\langle C\rangle\rangle \bigcirc \phi$): coalition C has a joint strategy to bring about ϕ.

M. Pauly. *A Modal Logic for Coalition Powers in Games. Journal of Logic and Computation* **12** (2002).

An Example

Two agents, A and B, must choose between two outcomes, p and q. We want a mechanism that will allow them to choose, which will satisfy the following requirements:

1. We definitely want an outcome to result, i.e., either p or q must be selected
2. We want the agents to be able to collectively choose and outcome
3. We do not want them to be able to bring about both outcomes simultaneously
4. We want them both to have equal power

An Example

Two agents, A and B, must choose between two outcomes, p and q. We want a mechanism that will allow them to choose, which will satisfy the following requirements:

1. We definitely want an outcome to result, i.e., either p or q must be selected: $[\emptyset](p \vee q)$
2. We want the agents to be able to collectively choose and outcome
3. We do not want them to be able to bring about both outcomes simultaneously
4. We want them both to have equal power

An Example

Two agents, A and B, must choose between two outcomes, p and q. We want a mechanism that will allow them to choose, which will satisfy the following requirements:

1. We definitely want an outcome to result, i.e., either p or q must be selected: $[\emptyset](p \vee q)$
2. We want the agents to be able to collectively choose and outcome: $[A, B]p \wedge [A, B]q$
3. We do not want them to be able to bring about both outcomes simultaneously
4. We want them both to have equal power

An Example

Two agents, A and B, must choose between two outcomes, p and q. We want a mechanism that will allow them to choose, which will satisfy the following requirements:

1. We definitely want an outcome to result, i.e., either p or q must be selected: $[\emptyset](p \vee q)$
2. We want the agents to be able to collectively choose and outcome: $[A, B]p \wedge [A, B]q$
3. We do not want them to be able to bring about both outcomes simultaneously: $\neg[A, B](p \wedge q)$
4. We want them both to have equal power

An Example

Two agents, A and B, must choose between two outcomes, p and q. We want a mechanism that will allow them to choose, which will satisfy the following requirements:

1. We definitely want an outcome to result, i.e., either p or q must be selected: $[\emptyset](p \vee q)$
2. We want the agents to be able to collectively choose and outcome: $[A, B]p \wedge [A, B]q$
3. We do not want them to be able to bring about both outcomes simultaneously: $\neg[A, B](p \wedge q)$
4. We want them both to have equal power: $\neg[x]p \wedge \neg[x]q$ where $x \in \{A, B\}$

An Example

Consider the following mechanism:

The two agents vote on the outcomes, i.e., they choose either p or q. If there is a consensus, then the consensus is selected; if there is no consensus, then an outcome p or q is selected non-deterministically.

Pauly and Wooldridge use the MOCHA model checking system to verify that the above procedure satisfies the previous specifications.

See, for example,

M. Pauly. *A Modal Logic for Coalition Powers in Games. Journal of Logic and Computation* **12** (2002).

Goranko and Jamroga. *Comparing Semantics of Logics fro Multi-Agent Systems.* See the website.

Conclusion

What can Logic do for Game Theory?

See `staff.science.uva.nl/~epacuit/caputLLI.html` for more information.

Thank you.