Logic and Artificial Intelligence

Lecture 1

Eric Pacuit

Currently Visiting the Center for Formal Epistemology, CMU

Center for Logic and Philosophy of Science
Tilburg University

ai.stanford.edu/∼epacuit
e.j.pacuit@uvt.nl

August 31, 2011

Dynamic
Epistemic/Doxastic
Logic

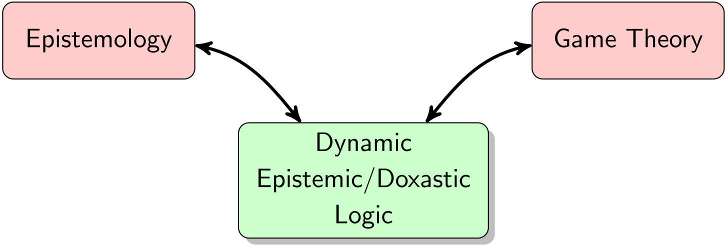
Epistemology
Game Theory
Dynamic Epistemic/Doxastic Logic

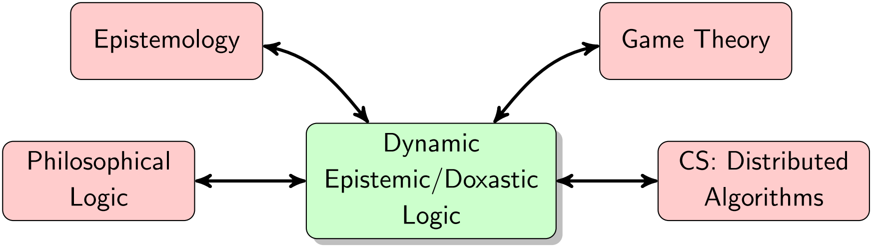
Epistemology
Game Theory
Philosophical Logic
Dynamic Epistemic/Doxastic Logic
CS: Distributed Algorithms

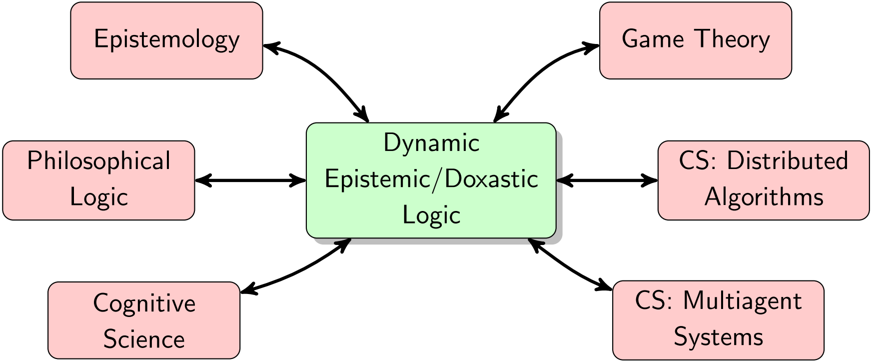
Epistemology
Game Theory
Philosophical Logic
Dynamic Epistemic/Doxastic Logic
CS: Distributed Algorithms
Cognitive Science
CS: Multiagent Systems

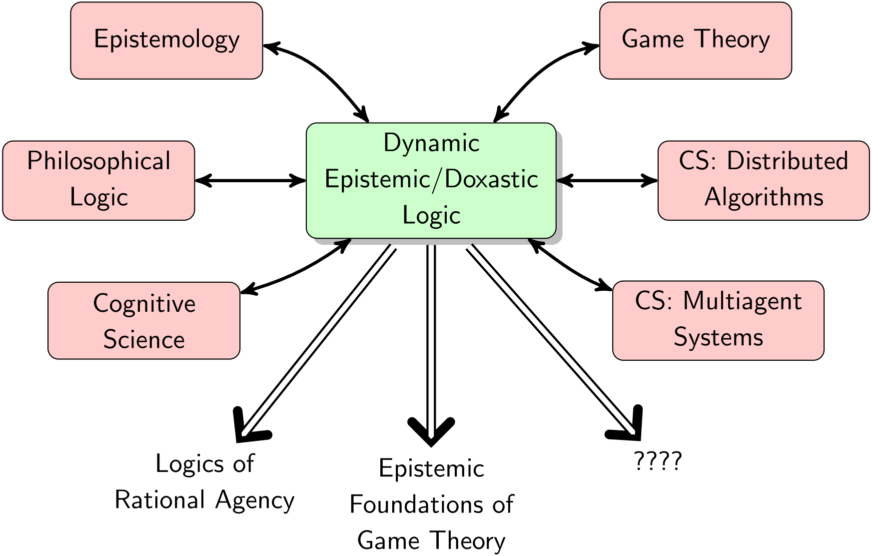
Epistemology
Game Theory
Philosophical Logic
Dynamic Epistemic/Doxastic Logic
CS: Distributed Algorithms
Cognitive Science
CS: Multiagent Systems
Logics of Rational Agency
Epistemic Foundations of Game Theory
????

Foundations of Epistemic Logic

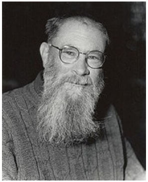

David Lewis

Jakko Hintikka

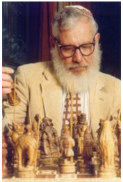

Robert Aumann

Larry Moss

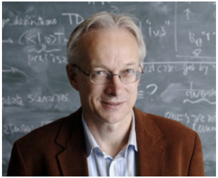

Johan van Benthem

Alexandru Baltag

Foundations of Epistemic Logic

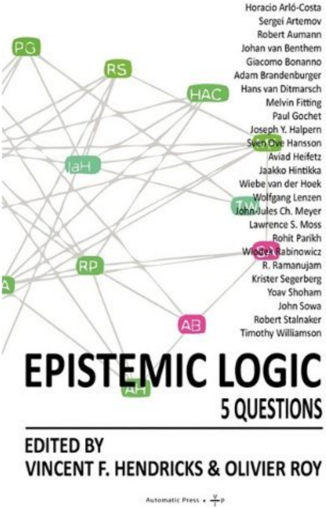

Single-Agent Epistemic Logic

Let KP informally mean "the agent knows that P (is true)".

Single-Agent Epistemic Logic

Let KP informally mean "the agent knows that P (is true)".

$K(P \to Q)$: "Ann knows that P implies Q"

Single-Agent Epistemic Logic

Let KP informally mean "the agent knows that P (is true)".

$K(P \to Q)$: "Ann knows that P implies Q"

$KP \lor \neg KP$: "either Ann does or does not know P"

Single-Agent Epistemic Logic

Let KP informally mean "the agent knows that P (is true)".

$K(P \to Q)$: "Ann knows that P implies Q"

$KP \lor \neg KP$: "either Ann does or does not know P"

$KP \lor K\neg P$: "Ann knows whether P is true"

Single-Agent Epistemic Logic

Let KP informally mean "the agent knows that P (is true)".

$K(P \to Q)$: "Ann knows that P implies Q"

$KP \vee \neg KP$: "either Ann does or does not know P"

$KP \vee K\neg P$: "Ann knows whether P is true"

LP: "P is an epistemic possibility"

Single-Agent Epistemic Logic

Let KP informally mean "the agent knows that P (is true)".

$K(P \to Q)$: "Ann knows that P implies Q"

$KP \vee \neg KP$: "either Ann does or does not know P"

$KP \vee K\neg P$: "Ann knows whether P is true"

LP: "P is an epistemic possibility"

KLP: "Ann knows that she thinks P is possible"

Example

Suppose there are three cards: 1, 2 and 3.

Ann is dealt one of the cards, one of the cards is placed face down on the table and the third card is put back in the deck.

Example

Suppose there are three cards: 1, 2 and 3.

Ann is dealt one of the cards, one of the cards is placed face down on the table and the third card is put back in the deck.

What are the relevant states?

Example

Suppose there are three cards: 1, 2 and 3.

Ann is dealt one of the cards, one of the cards is placed face down on the table and the third card is put back in the deck.

What are the relevant states?

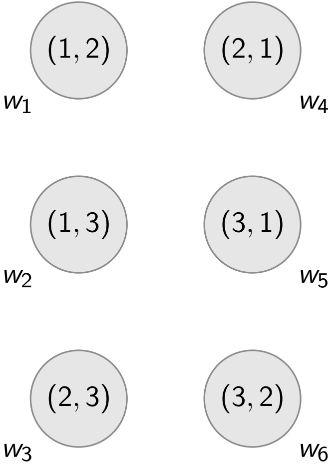

Example

Suppose there are three cards: 1, 2 and 3.

Ann is dealt one of the cards, one of the cards is placed face down on the table and the third card is put back in the deck.

Ann receives card 3 and card 1 is put on the table

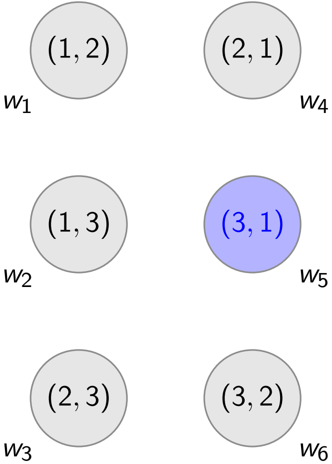

Example

Suppose there are three cards: 1, 2 and 3.

Ann is dealt one of the cards, one of the cards is placed face down on the table and the third card is put back in the deck.

What information does Ann have?

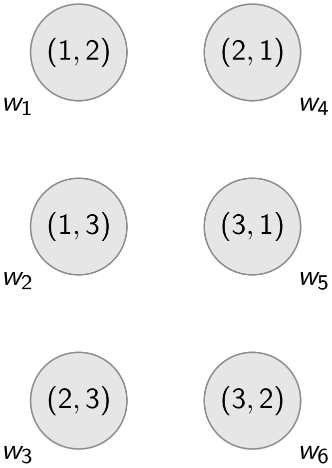

Example

Suppose there are three cards: 1, 2 and 3.

Ann is dealt one of the cards, one of the cards is placed face down on the table and the third card is put back in the deck.

What information does Ann have?

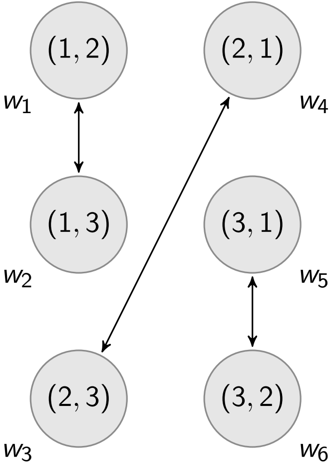

Example

Suppose there are three cards: 1, 2 and 3.

Ann is dealt one of the cards, one of the cards is placed face down on the table and the third card is put back in the deck.

What information does Ann have?

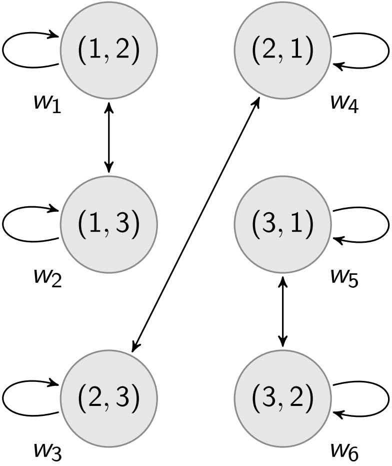

Example

Suppose there are three cards: 1, 2 and 3.

Ann is dealt one of the cards, one of the cards is placed face down on the table and the third card is put back in the deck.

Suppose H_i is intended to mean "Ann has card i"

T_i is intended to mean "card i is on the table"

Eg., $V(H_1) = \{w_1, w_2\}$

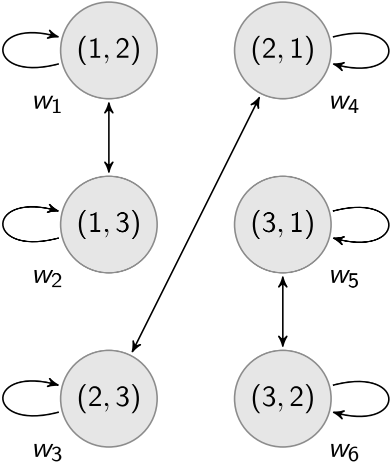

Example

Suppose there are three cards: 1, 2 and 3.

Ann is dealt one of the cards, one of the cards is placed face down on the table and the third card is put back in the deck.

Suppose H_i is intended to mean "Ann has card i"

T_i is intended to mean "card i is on the table"

Eg., $V(H_1) = \{w_1, w_2\}$

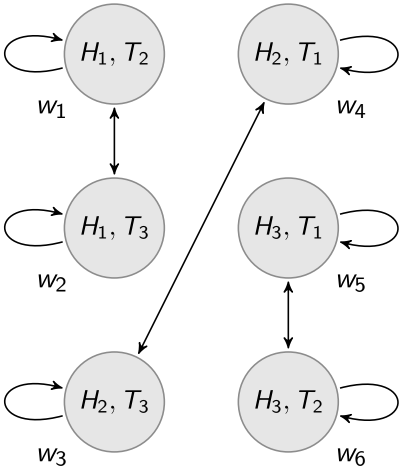

Example

Suppose there are three cards: 1, 2 and 3.

Ann is dealt one of the cards, one of the cards is placed face down on the table and the third card is put back in the deck.

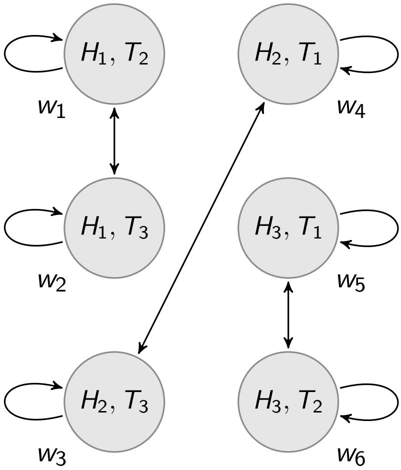

Example

Suppose there are three cards: 1, 2 and 3.

Ann is dealt one of the cards, one of the cards is placed face down on the table and the third card is put back in the deck.

Suppose that Ann receives card 1 and card 2 is on the table.

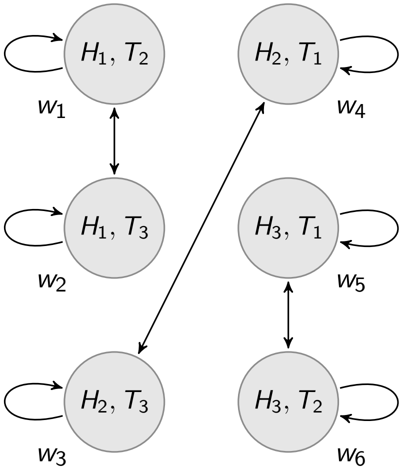

Example

Suppose there are three cards: 1, 2 and 3.

Ann is dealt one of the cards, one of the cards is placed face down on the table and the third card is put back in the deck.

Suppose that Ann receives card 1 and card 2 is on the table.

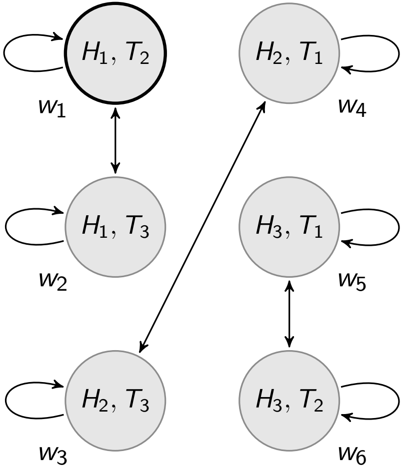

Example

Suppose there are three cards: 1, 2 and 3.

Ann is dealt one of the cards, one of the cards is placed face down on the table and the third card is put back in the deck.

$\mathcal{M}, w_1 \models KH_1$

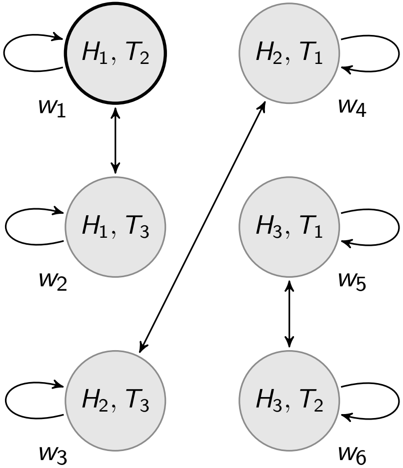

Example

Suppose there are three cards: 1, 2 and 3.

Ann is dealt one of the cards, one of the cards is placed face down on the table and the third card is put back in the deck.

$\mathcal{M}, w_1 \models K H_1$

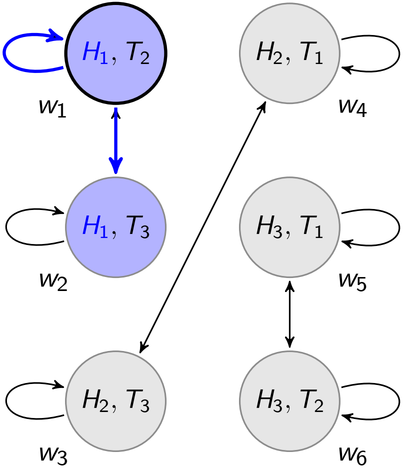

Example

Suppose there are three cards: 1, 2 and 3.

Ann is dealt one of the cards, one of the cards is placed face down on the table and the third card is put back in the deck.

$\mathcal{M}, w_1 \models K H_1$

$\mathcal{M}, w_1 \models K \neg T_1$

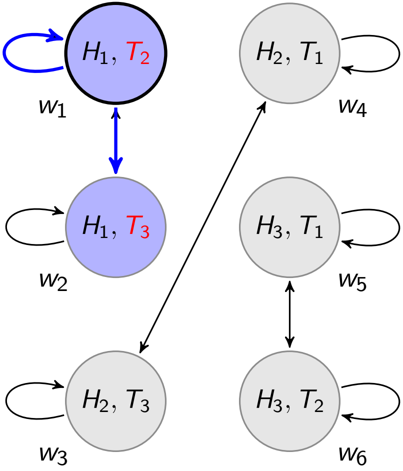

Example

Suppose there are three cards: 1, 2 and 3.

Ann is dealt one of the cards, one of the cards is placed face down on the table and the third card is put back in the deck.

$\mathcal{M}, w_1 \models LT_2$

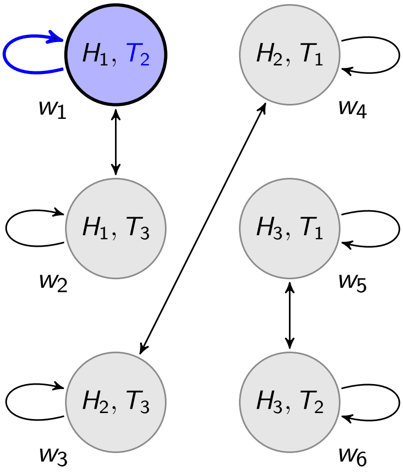

Example

Suppose there are three cards: 1, 2 and 3.

Ann is dealt one of the cards, one of the cards is placed face down on the table and the third card is put back in the deck.

$\mathcal{M}, w_1 \models K(T_2 \vee T_3)$

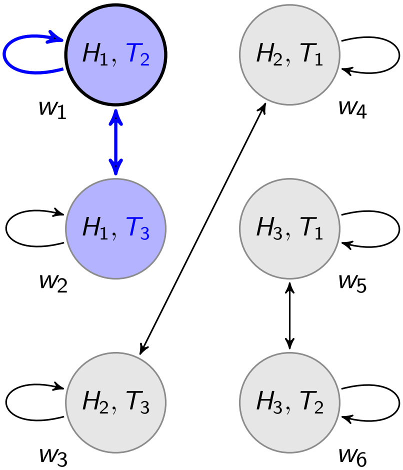

Multiagent Epistemic Logic

Many of the examples we are interested in involve more than one agent!

Multiagent Epistemic Logic

Many of the examples we are interested in involve more than one agent!

$K_A P$ means "Ann knows P"

$K_B P$ means "Bob knows P"

Multiagent Epistemic Logic

Many of the examples we are interested in involve more than one agent!

$K_A P$ means "Ann knows P"

$K_B P$ means "Bob knows P"

- $K_A K_B \varphi$: "Ann knows that Bob knows φ"
- $K_A(K_B\varphi \vee K_B\neg\varphi)$: "Ann knows that Bob knows whether φ
- $\neg K_B K_A K_B(\varphi)$: "Bob does not know that Ann knows that Bob knows that φ"

Example

Suppose there are three cards: 1, 2 and 3.

Ann is dealt one of the cards, one of the cards is placed face down on the table and the third card is put back in the deck.

Suppose that Ann receives card 1 and card 2 is on the table.

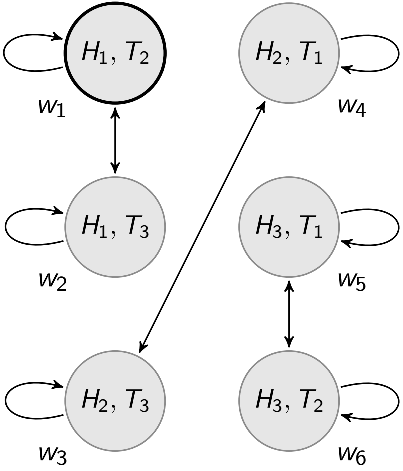

Example

Suppose there are three cards: 1, 2 and 3.

Ann is dealt one of the cards, Bob is given one of the cards and the third card is put back in the deck.

Suppose that Ann receives card 1 and Bob receives card 2.

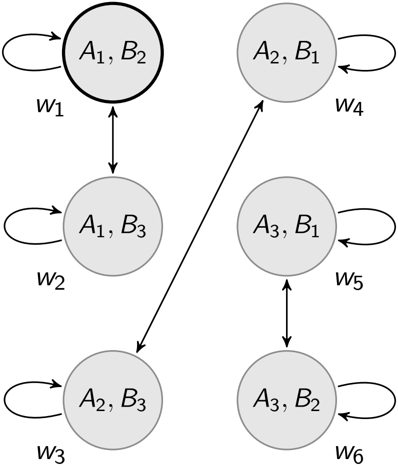

Example

Suppose there are three cards: 1, 2 and 3.

Ann is dealt one of the cards, Bob is given one of the cards and the third card is put back in the deck.

Suppose that Ann receives card 1 and Bob receives card 2.

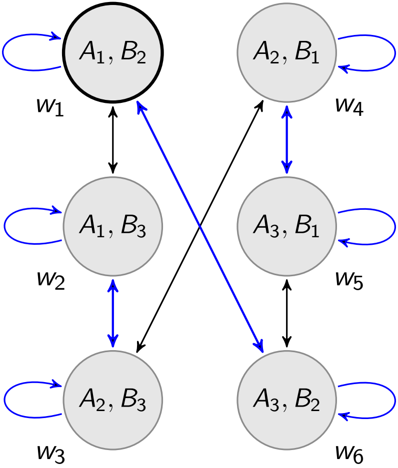

Example

Suppose there are three cards: 1, 2 and 3.

Ann is dealt one of the cards, Bob is given one of the cards and the third card is put back in the deck.

Suppose that Ann receives card 1 and Bob receives card 2.

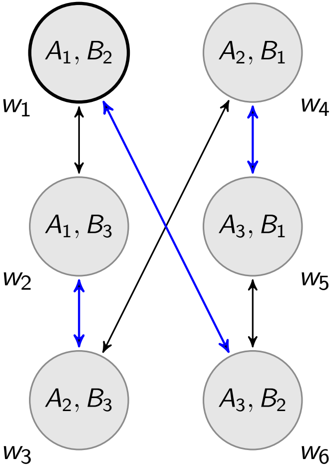

Example

Suppose there are three cards: 1, 2 and 3.

Ann is dealt one of the cards, Bob is given one of the cards and the third card is put back in the deck.

Suppose that Ann receives card 1 and Bob receives card 2.

$\mathcal{M}, w_1 \models K_B(K_A A_1 \vee K_A \neg A_1)$

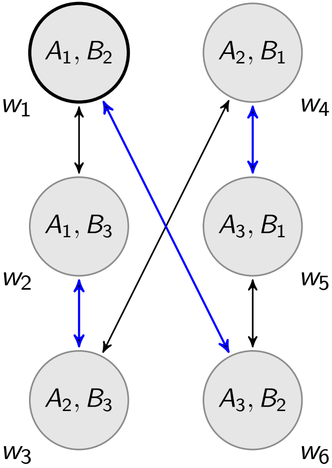

Example

Suppose there are three cards: 1, 2 and 3.

Ann is dealt one of the cards, Bob is given one of the cards and the third card is put back in the deck.

Suppose that Ann receives card 1 and Bob receives card 2.

$\mathcal{M}, w_1 \models K_B(K_A A_1 \vee K_A \neg A_1)$

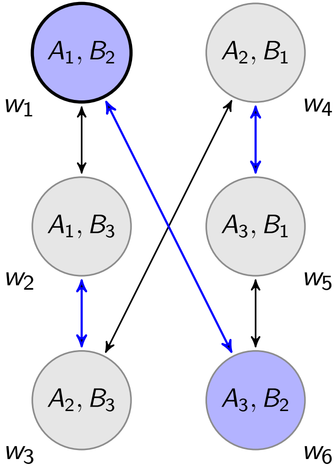

Example

Suppose there are three cards: 1, 2 and 3.

Ann is dealt one of the cards, Bob is given one of the cards and the third card is put back in the deck.

Suppose that Ann receives card 1 and Bob receives card 2.

$\mathcal{M}, w_1 \models K_B(K_A A_1 \vee K_A \neg A_1)$

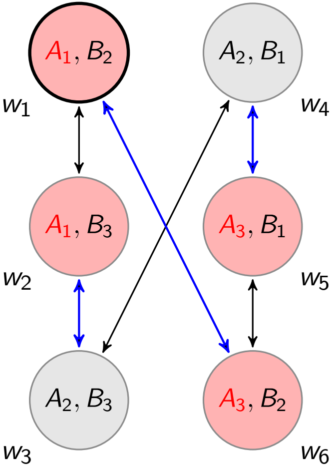

Three children are outside playing. Two of them get mud on their forehead. They cannot see or feel the mud on their own foreheads, but can see who is dirty.

Three children are outside playing. Two of them get mud on their forehead. They cannot see or feel the mud on their own foreheads, but can see who is dirty.

Their mother enters the room and says "At least one of you have mud on your forehead".

Three children are outside playing. Two of them get mud on their forehead. They cannot see or feel the mud on their own foreheads, but can see who is dirty.

Their mother enters the room and says "At least one of you have mud on your forehead".

Then the children are repeatedly asked "do you know if you have mud on your forehead?"

Three children are outside playing. Two of them get mud on their forehead. They cannot see or feel the mud on their own foreheads, but can see who is dirty.

Their mother enters the room and says "At least one of you have mud on your forehead".

Then the children are repeatedly asked "do you know if you have mud on your forehead?"

What happens?

Three children are outside playing. Two of them get mud on their forehead. They cannot see or feel the mud on their own foreheads, but can see who is dirty.

Their mother enters the room and says "At least one of you have mud on your forehead".

Then the children are repeatedly asked "do you know if you have mud on your forehead?"

What happens?

Claim: After first question, the children answer "I don't know",

Three children are outside playing. Two of them get mud on their forehead. They cannot see or feel the mud on their own foreheads, but can see who is dirty.

Their mother enters the room and says "At least one of you have mud on your forehead".

Then the children are repeatedly asked "do you know if you have mud on your forehead?"

What happens?

Claim: After first question, the children answer "I don't know", after the second question the muddy children answer "I have mud on my forehead!" (but the clean child is still in the dark).

Three children are outside playing. Two of them get mud on their forehead. They cannot see or feel the mud on their own foreheads, but can see who is dirty.

Their mother enters the room and says "At least one of you have mud on your forehead".

Then the children are repeatedly asked "do you know if you have mud on your forehead?"

What happens?

Claim: After first question, the children answer "I don't know", after the second question the muddy children answer "I have mud on my forehead!" (but the clean child is still in the dark). Then the clean child says, "Oh, I must be clean."

Three children are outside playing. Two of them get mud on their forehead. They cannot see or feel the mud on their own foreheads, but can see who is dirty.

Their mother enters the room and says “At least one of you have mud on your forehead”.

Then the children are repeatedly asked “do you know if you have mud on your forehead?”

What happens?

Claim: After first question, the children answer “I don’t know”, after the second question the muddy children answer “I have mud on my forehead!” (but the clean child is still in the dark). Then the clean child says, “Oh, I must be clean.” ‣ Please, I beg you not to go over this!

Muddy Children

Assume:

- There are three children: Ann, Bob and Charles.
- (Only) Ann and Bob have mud on their forehead.

Muddy Children

Assume:

- There are three children: Ann, Bob and Charles.
- (Only) Ann and Bob have mud on their forehead.

Muddy Children

Assume:

- There are three children: Ann, Bob and Charles.
- (Only) Ann and Bob have mud on their forehead.

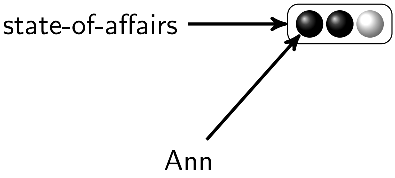

Muddy Children

Assume:

- There are three children: Ann, Bob and Charles.
- (Only) Ann and Bob have mud on their forehead.

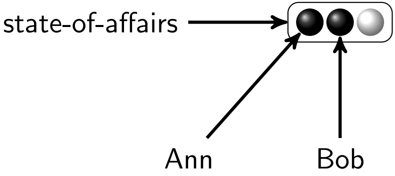

Muddy Children

Assume:

- There are three children: Ann, Bob and Charles.
- (Only) Ann and Bob have mud on their forehead.

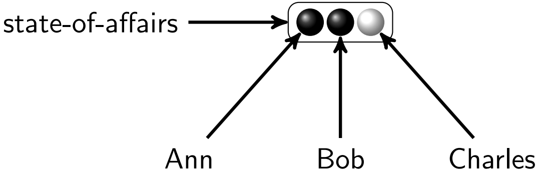

Muddy Children

Assume:

- There are three children: Ann, Bob and Charles.
- (Only) Ann and Bob have mud on their forehead.

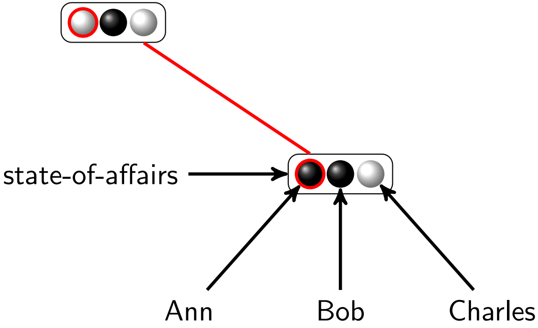

Muddy Children

Assume:

- There are three children: Ann, Bob and Charles.
- (Only) Ann and Bob have mud on their forehead.

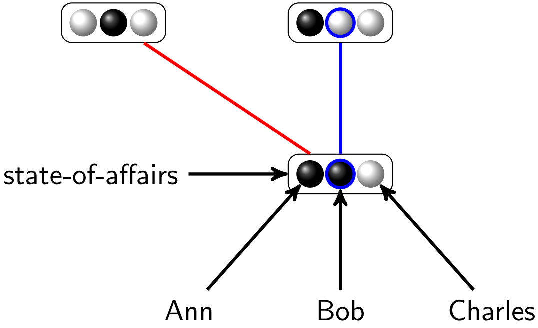

Muddy Children

Assume:

- There are three children: Ann, Bob and Charles.
- (Only) Ann and Bob have mud on their forehead.

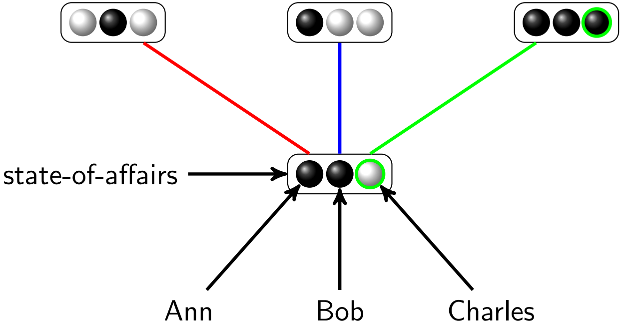

Muddy Children

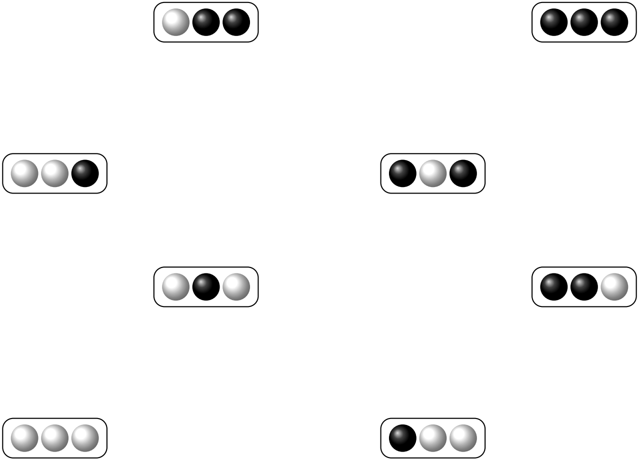

All 8 possible situations

Muddy Children

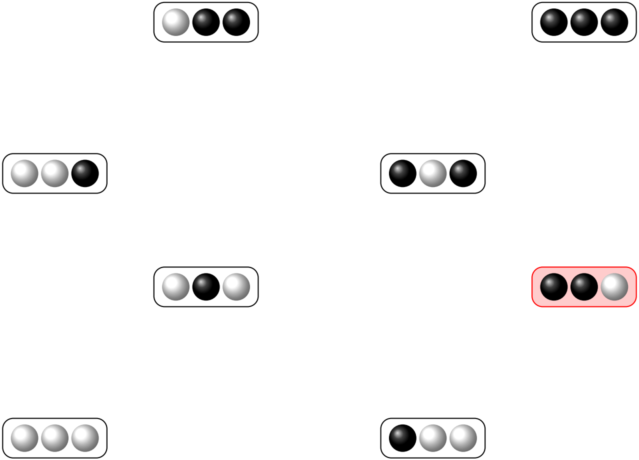

The actual situation

Muddy Children

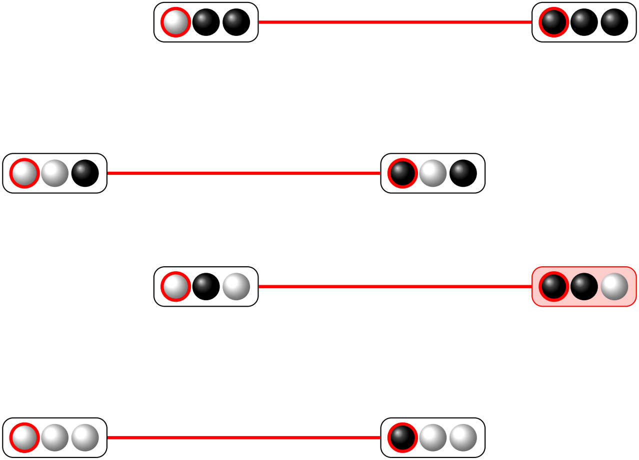

Ann's uncertainty

Muddy Children

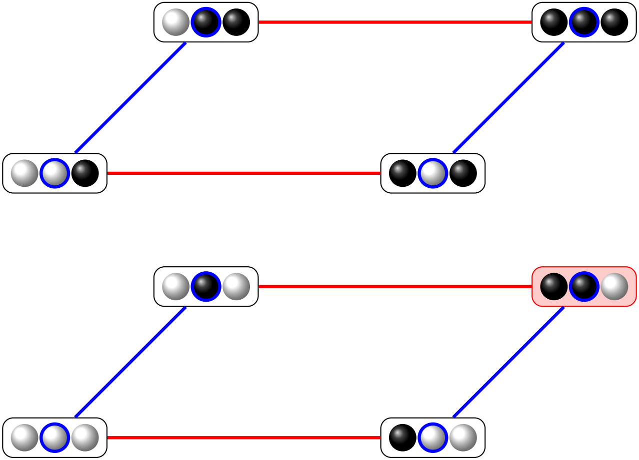

Bob's uncertainty

Muddy Children

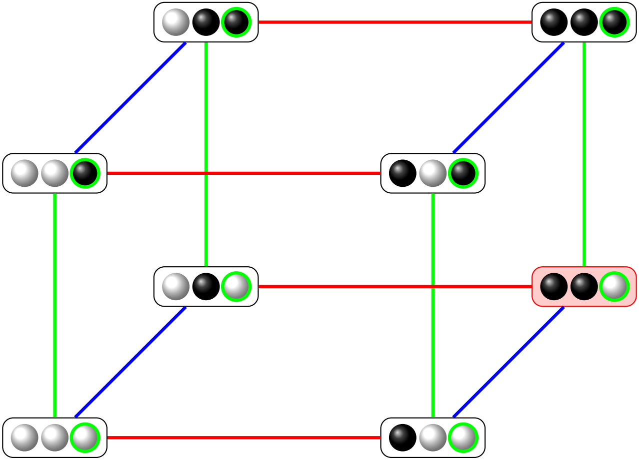

Charles' uncertainty

Muddy Children

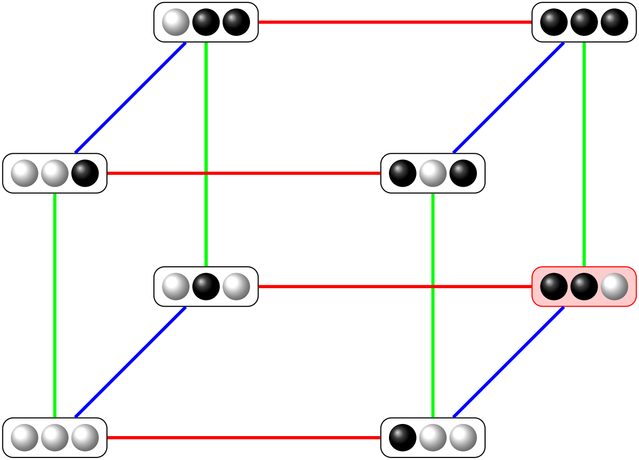

Muddy Children

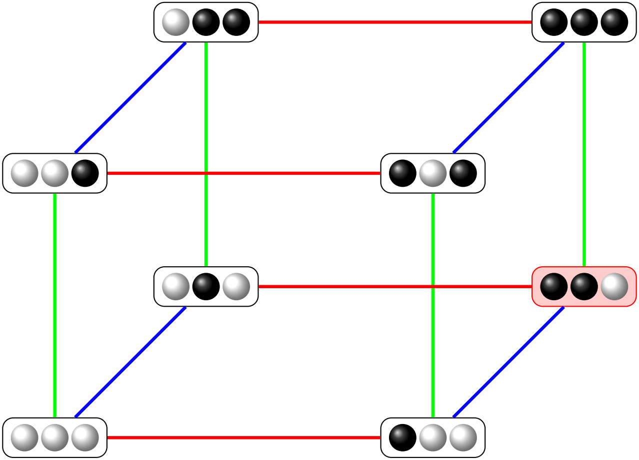

None of the children know if they are muddy

Muddy Children

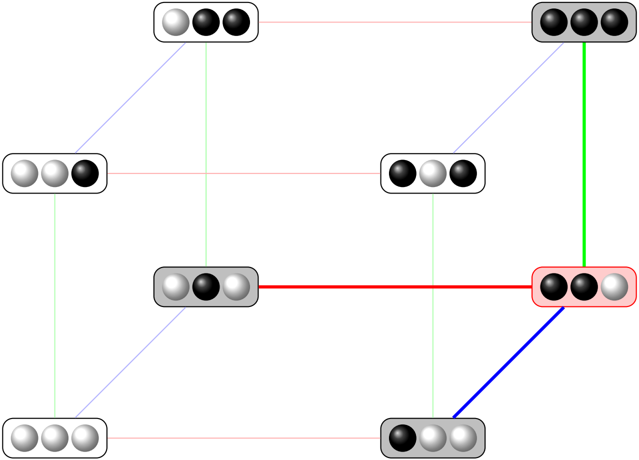

None of the children know if they are muddy

Muddy Children

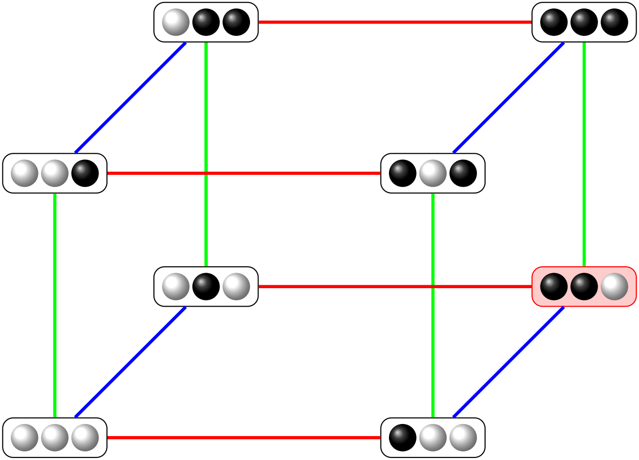

"At least one has mud on their forehead."

Muddy Children

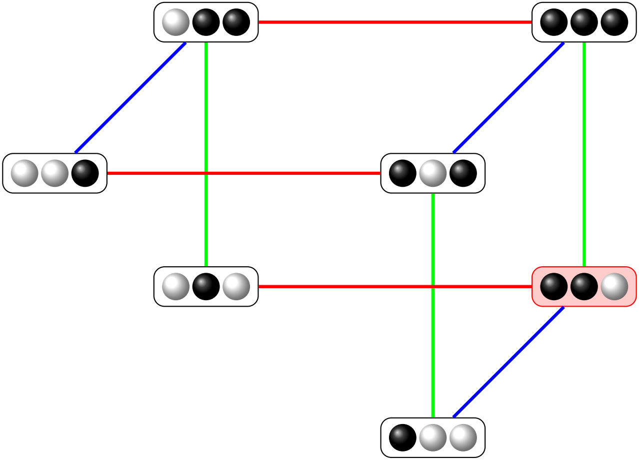

"At least one has mud on their forehead."

Muddy Children

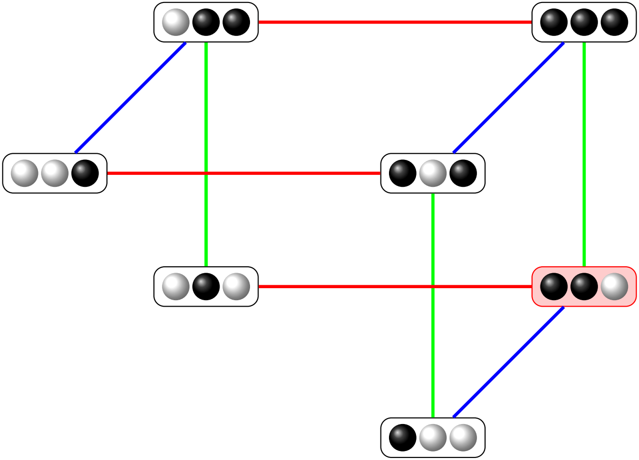

"Who has mud on their forehead?"

Muddy Children

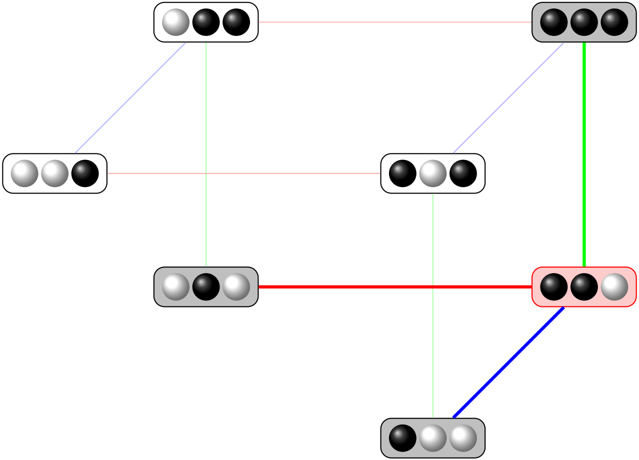

"Who has mud on their forehead?"

Muddy Children

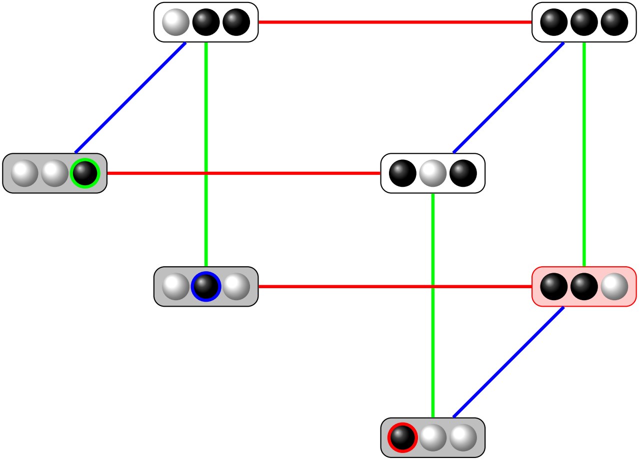

No one steps forward.

Muddy Children

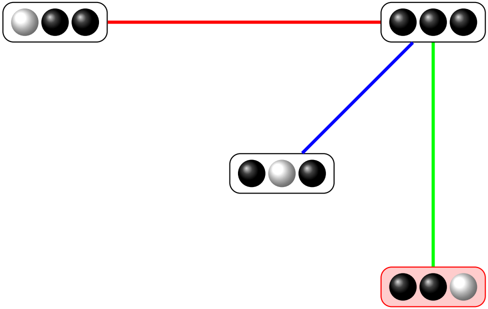

No one steps forward.

Muddy Children

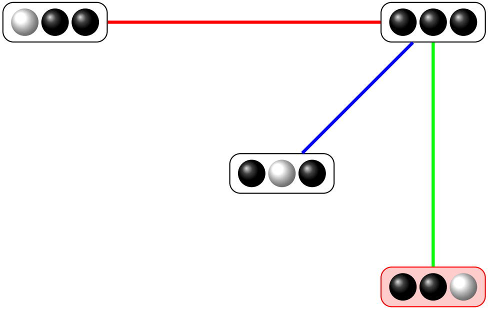

“Who has mud on their forehead?”

Muddy Children

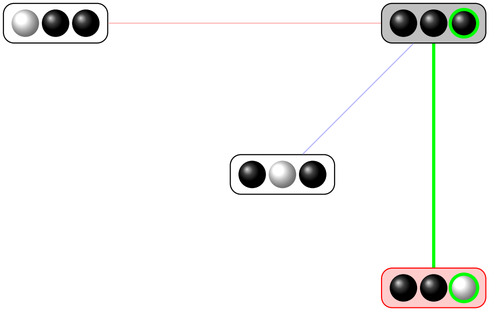

Charles does not know he is clean.

Muddy Children

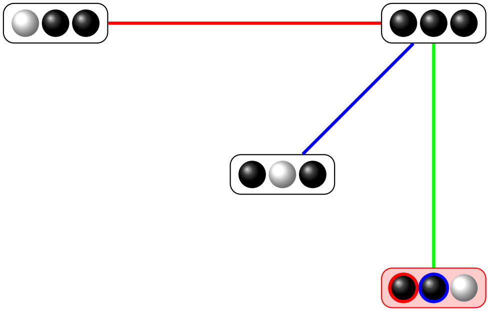

Ann and Bob step forward.

Muddy Children

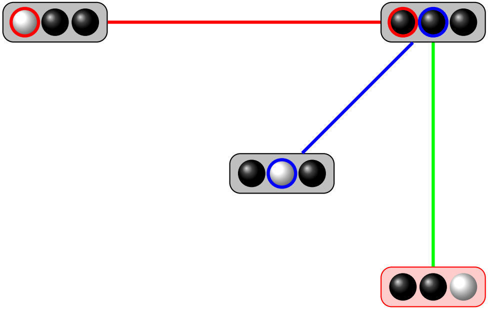

Now, Charles knows he is clean.

Muddy Children

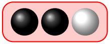

Now, Charles knows he is clean.

Single-Agent Epistemic Logic: The Language

φ is a formula of Epistemic Logic ($\mathcal{L}$) if it is of the form

$$\varphi \;:=\; p \mid \neg\varphi \mid \varphi \wedge \psi \mid K\varphi$$

Single-Agent Epistemic Logic: The Language

φ is a formula of Epistemic Logic ($\mathcal{L}$) if it is of the form

$$\varphi \ := \ p \mid \neg\varphi \mid \varphi \wedge \psi \mid K\varphi$$

- $p \in \mathsf{At}$ is an atomic fact.
  - "It is raining"
  - "The talk is at 2PM"
  - "The card on the table is a 7 of Hearts"

Single-Agent Epistemic Logic: The Language

φ is a formula of Epistemic Logic ($\mathcal{L}$) if it is of the form

$$\varphi \ := \ p \mid \neg\varphi \mid \varphi \wedge \psi \mid K\varphi$$

- $p \in \mathsf{At}$ is an atomic fact.
- The usual propositional language ($\mathcal{L}_0$)

Single-Agent Epistemic Logic: The Language

φ is a formula of Epistemic Logic ($\mathcal{L}$) if it is of the form

$$\varphi \ := \ p \mid \neg\varphi \mid \varphi \wedge \psi \mid K\varphi$$

- $p \in \mathsf{At}$ is an atomic fact.
- The usual propositional language ($\mathcal{L}_0$)
- $K\varphi$ is intended to mean "The agent knows that φ is true".

Single-Agent Epistemic Logic: The Language

φ is a formula of Epistemic Logic ($\mathcal{L}$) if it is of the form

$$\varphi \;:=\; p \mid \neg\varphi \mid \varphi \wedge \psi \mid K\varphi$$

- $p \in \mathsf{At}$ is an atomic fact.
- The usual propositional language ($\mathcal{L}_0$)
- $K\varphi$ is intended to mean "The agent knows that φ is true".
- The usual definitions for $\rightarrow, \vee, \leftrightarrow$ apply
- Define $L\varphi$ as $\neg K \neg\varphi$

Single-Agent Epistemic Logic: The Language

φ is a formula of Epistemic Logic ($\mathcal{L}$) if it is of the form

$$\varphi \; := \; p \mid \neg\varphi \mid \varphi \wedge \psi \mid K\varphi$$

$K(p \to q)$: "Ann knows that p implies q"

$Kp \vee \neg Kp$:

$Kp \vee K\neg p$:

$L\varphi$:

$KL\varphi$:

Single-Agent Epistemic Logic: The Language

φ is a formula of Epistemic Logic ($\mathcal{L}$) if it is of the form

$$\varphi \ := \ p \mid \neg\varphi \mid \varphi \wedge \psi \mid K\varphi$$

$K(p \to q)$: "Ann knows that p implies q"

$Kp \vee \neg Kp$: "either Ann does or does not know p"

$Kp \vee K\neg p$: "Ann knows whether p is true"

$L\varphi$:

$KL\varphi$:

Single-Agent Epistemic Logic: The Language

φ is a formula of Epistemic Logic ($\mathcal{L}$) if it is of the form

$$\varphi \ := \ p \mid \neg\varphi \mid \varphi \wedge \psi \mid K\varphi$$

$K(p \rightarrow q)$: "Ann knows that p implies q"

$Kp \vee \neg Kp$: "either Ann does or does not know p"

$Kp \vee K\neg p$: "Ann knows whether p is true"

$L\varphi$: "φ is an epistemic possibility"

$KL\varphi$: "Ann knows that she thinks φ is possible"

Single-Agent Epistemic Logic: Kripke Models

$$\mathcal{M} = \langle W, R, V \rangle$$

Single-Agent Epistemic Logic: Kripke Models

$$\mathcal{M} = \langle W, R, V \rangle$$

- $W \neq \emptyset$ is the set of all relevant situations (states of affairs, possible worlds)

Single-Agent Epistemic Logic: Kripke Models

$$\mathcal{M} = \langle W, R, V \rangle$$

- $W \neq \emptyset$ is the set of all relevant situations (states of affairs, possible worlds)
- $R \subseteq W \times W$ represents the information of the agent: wRv provided "w and v are epistemically indistinguishable"

Single-Agent Epistemic Logic: Kripke Models

$$\mathcal{M} = \langle W, R, V \rangle$$

- $W \neq \emptyset$ is the set of all relevant situations (states of affairs, possible worlds)
- $R \subseteq W \times W$ represents the information of the agent:
  wRv provided "w and v are epistemically indistinguishable"
  Notation: many people write $\sim$ for R
- $V : \mathsf{At} \to \wp(W)$ is a valuation function assigning propositional variables to worlds

Single Agent Epistemic Logic: Truth in a Model

Given $\varphi \in \mathcal{L}$, a Kripke model $\mathcal{M} = \langle W, R, V \rangle$ and $w \in W$

$\mathcal{M}, w \models \varphi$ means "in $\mathcal{M}$, if the actual state is w, then φ is true"

Single Agent Epistemic Logic: Truth in a Model

Given $\varphi \in \mathcal{L}$, a Kripke model $\mathcal{M} = \langle W, R, V \rangle$ and $w \in W$

$\mathcal{M}, w \models \varphi$ is defined as follows:

- $\mathcal{M}, w \models p$ iff $w \in V(p)$ (with $p \in \mathrm{At}$)
- $\mathcal{M}, w \models \neg\varphi$ if $\mathcal{M}, w \not\models \varphi$
- $\mathcal{M}, w \models \varphi \wedge \psi$ if $\mathcal{M}, w \models \varphi$ and $\mathcal{M}, w \models \psi$
- $\mathcal{M}, w \models K\varphi$ if for each $v \in W$, if wRv, then $\mathcal{M}, v \models \varphi$

Single Agent Epistemic Logic: Truth in a Model

Given $\varphi \in \mathcal{L}$, a Kripke model $\mathcal{M} = \langle W, R, V \rangle$ and $w \in W$

$\mathcal{M}, w \models \varphi$ is defined as follows:

- ✓ $\mathcal{M}, w \models p$ iff $w \in V(p)$ (with $p \in \text{At}$)
- $\mathcal{M}, w \models \neg\varphi$ if $\mathcal{M}, w \not\models \varphi$
- $\mathcal{M}, w \models \varphi \wedge \psi$ if $\mathcal{M}, w \models \varphi$ and $\mathcal{M}, w \models \psi$
- $\mathcal{M}, w \models K\varphi$ if for each $v \in W$, if wRv, then $\mathcal{M}, v \models \varphi$

Single Agent Epistemic Logic: Truth in a Model

Given $\varphi \in \mathcal{L}$, a Kripke model $\mathcal{M} = \langle W, R, V \rangle$ and $w \in W$

$\mathcal{M}, w \models \varphi$ is defined as follows:

- ✓ $\mathcal{M}, w \models p$ iff $w \in V(p)$ (with $p \in \mathrm{At}$)
- ✓ $\mathcal{M}, w \models \neg\varphi$ if $\mathcal{M}, w \not\models \varphi$
- ✓ $\mathcal{M}, w \models \varphi \wedge \psi$ if $\mathcal{M}, w \models \varphi$ and $\mathcal{M}, w \models \psi$
- $\mathcal{M}, w \models K\varphi$ if for each $v \in W$, if wRv, then $\mathcal{M}, v \models \varphi$

Single Agent Epistemic Logic: Truth in a Model

Given $\varphi \in \mathcal{L}$, a Kripke model $\mathcal{M} = \langle W, R, V \rangle$ and $w \in W$

$\mathcal{M}, w \models \varphi$ is defined as follows:

- ✓ $\mathcal{M}, w \models p$ iff $w \in V(p)$ (with $p \in \mathrm{At}$)
- ✓ $\mathcal{M}, w \models \neg\varphi$ if $\mathcal{M}, w \not\models \varphi$
- ✓ $\mathcal{M}, w \models \varphi \wedge \psi$ if $\mathcal{M}, w \models \varphi$ and $\mathcal{M}, w \models \psi$
- ✓ $\mathcal{M}, w \models K\varphi$ if for each $v \in W$, if wRv, then $\mathcal{M}, v \models \varphi$

Some Questions

Should we make additional assumptions about R (i.e., reflexive, transitive, etc.)

What idealizations have we made?

Some Notation

A Kripke Frame is a tuple $\langle W, R \rangle$ where $R \subseteq W \times W$.

φ is valid in a Kripke model $\mathcal{M}$ if $\mathcal{M}, w \models \varphi$ for all states w (we write $\mathcal{M} \models \varphi$).

φ is valid on a Kripke frame $\mathcal{F}$ if $\mathcal{M} \models \varphi$ for all models $\mathcal{M}$ based on $\mathcal{F}$.

Logical Omniscience

Fact: φ is valid then $K\varphi$ is valid

Logical Omniscience

Fact: $K\varphi \wedge K\psi \to K(\varphi \wedge \psi)$ is valid on all Kripke frames

Logical Omniscience

Fact: If $\varphi \to \psi$ is valid then $K\varphi \to K\psi$ is valid

Logical Omniscience

Fact: $K(\varphi \to \psi) \to (K\varphi \to K\psi)$ is valid on all Kripke frames.

Logical Omniscience

Fact: $\varphi \leftrightarrow \psi$ is valid then $K\varphi \leftrightarrow K\psi$ is valid

Correspondence

Definition

A model formula φ corresponds to a property P (of a relation in a Kripke frame) provided

$$\mathcal{F} \models \varphi \text{ iff } \mathcal{F} \text{ has } P$$

Correspondence

Definition

A model formula φ corresponds to a property P (of a relation in a Kripke frame) provided

$$\mathcal{F} \models \varphi \text{ iff } \mathcal{F} \text{ has } P$$

Modal Formula	Corresponding Property

Correspondence

Definition

A model formula φ corresponds to a property P (of a relation in a Kripke frame) provided

$$\mathcal{F} \models \varphi \text{ iff } \mathcal{F} \text{ has } P$$

Modal Formula	Corresponding Property
$K\varphi \to \varphi$	Reflexive

Correspondence

Definition

A model formula φ corresponds to a property P (of a relation in a Kripke frame) provided

$$\mathcal{F} \models \varphi \text{ iff } \mathcal{F} \text{ has } P$$

Modal Formula	Corresponding Property
$K\varphi \to \varphi$	Reflexive
$K\varphi \to KK\varphi$	Transitive

Correspondence

Definition

A model formula φ corresponds to a property P (of a relation in a Kripke frame) provided

$$\mathcal{F} \models \varphi \text{ iff } \mathcal{F} \text{ has } P$$

Modal Formula	Corresponding Property
$K\varphi \rightarrow \varphi$	Reflexive
$K\varphi \rightarrow KK\varphi$	Transitive
$\neg K\varphi \rightarrow K\neg K\varphi$	Euclidean

Correspondence

Definition

A model formula φ corresponds to a property P (of a relation in a Kripke frame) provided

$$\mathcal{F} \models \varphi \text{ iff } \mathcal{F} \text{ has } P$$

Modal Formula	Corresponding Property
$K\varphi \to \varphi$	Reflexive
$K\varphi \to KK\varphi$	Transitive
$\neg K\varphi \to K\neg K\varphi$	Euclidean
$\varphi \to KL\varphi$	Symmetric

Correspondence

Definition

A model formula φ corresponds to a property P (of a relation in a Kripke frame) provided

$$\mathcal{F} \models \varphi \text{ iff } \mathcal{F} \text{ has } P$$

Modal Formula	Corresponding Property
$K\varphi \to \varphi$	Reflexive
$K\varphi \to KK\varphi$	Transitive
$\neg K\varphi \to K\neg K\varphi$	Euclidean
$\varphi \to KL\varphi$	Symmetric
$\neg K\bot$	Serial

Modal Formula	Property	Philosophical Assumption

Modal Formula	Property	Philosophical Assumption
$K(\varphi \to \psi) \to (K\varphi \to K\psi)$	—	Logical Omniscience

Modal Formula	Property	Philosophical Assumption
$K(\varphi \to \psi) \to (K\varphi \to K\psi)$	—	Logical Omniscience
$K\varphi \to \varphi$	Reflexive	Truth

Modal Formula	Property	Philosophical Assumption
$K(\varphi \to \psi) \to (K\varphi \to K\psi)$	—	Logical Omniscience
$K\varphi \to \varphi$	Reflexive	Truth
$K\varphi \to KK\varphi$	Transitive	Positive Introspection

Modal Formula	Property	Philosophical Assumption
$K(\varphi \to \psi) \to (K\varphi \to K\psi)$	—	Logical Omniscience
$K\varphi \to \varphi$	Reflexive	Truth
$K\varphi \to KK\varphi$	Transitive	Positive Introspection
$\neg K\varphi \to K\neg K\varphi$	Euclidean	Negative Introspection

Modal Formula	Property	Philosophical Assumption
$K(\varphi \to \psi) \to (K\varphi \to K\psi)$	—	Logical Omniscience
$K\varphi \to \varphi$	Reflexive	Truth
$K\varphi \to KK\varphi$	Transitive	Positive Introspection
$\neg K\varphi \to K\neg K\varphi$	Euclidean	Negative Introspection
$\neg K\bot$	Serial	Consistency

The Logic **S5**

The logic **S5** contains the following axiom schemes and rules:

Pc	Axiomatization of Propositional Calculus
K	$K(\varphi \to \psi) \to (K\varphi \to K\psi)$
T	$K\varphi \to \varphi$
4	$K\varphi \to KK\varphi$
5	$\neg K\varphi \to K\neg K\varphi$
MP	$\dfrac{\varphi \qquad \varphi \to \psi}{\psi}$
Nec	$\dfrac{\varphi}{K\psi}$

The Logic **S5**

The logic **S5** contains the following axiom schemes and rules:

Pc	Axiomatization of Propositional Calculus
K	$K(\varphi \to \psi) \to (K\varphi \to K\psi)$
T	$K\varphi \to \varphi$
4	$K\varphi \to KK\varphi$
5	$\neg K\varphi \to K\neg K\varphi$
MP	$\dfrac{\varphi \qquad \varphi \to \psi}{\psi}$
Nec	$\dfrac{\varphi}{K\psi}$

Theorem

S5 *is sound and strongly complete with respect to the class of Kripke frames with equivalence relations.*

The Logic **S5**

The logic **S5** contains the following axiom schemes and rules:

Pc	Axiomatization of Propositional Calculus
K	$K(\varphi \to \psi) \to (K\varphi \to K\psi)$
T	$K\varphi \to \varphi$
4	$K\varphi \to KK\varphi$
5	$\neg K\varphi \to K\neg K\varphi$
MP	$\dfrac{\varphi \qquad \varphi \to \psi}{\psi}$
Nec	$\dfrac{\varphi}{K\psi}$

Theorem

S5 *is sound and strongly complete with respect to the class of Kripke frames with equivalence relations.*

The Logic **S5**

The logic **S5** contains the following axiom schemes and rules:

Pc	Axiomatization of Propositional Calculus
K	$K(\varphi \to \psi) \to (K\varphi \to K\psi)$
T	$K\varphi \to \varphi$
4	$K\varphi \to KK\varphi$
5	$\neg K\varphi \to K\neg K\varphi$
MP	$\dfrac{\varphi \qquad \varphi \to \psi}{\psi}$
Nec	$\dfrac{\varphi}{K\psi}$

Theorem

S5 *is sound and strongly complete with respect to the class of Kripke frames with equivalence relations.*

The Logic **S5**

The logic **S5** contains the following axiom schemes and rules:

Pc	Axiomatization of Propositional Calculus
K	$K(\varphi \to \psi) \to (K\varphi \to K\psi)$
T	$K\varphi \to \varphi$
4	$K\varphi \to KK\varphi$
5	$\neg K\varphi \to K\neg K\varphi$
MP	$\dfrac{\varphi \qquad \varphi \to \psi}{\psi}$
Nec	$\dfrac{\varphi}{K\psi}$

Theorem

S5 *is sound and strongly complete with respect to the class of Kripke frames with equivalence relations.*

The Logic **S5**

The logic **S5** contains the following axiom schemes and rules:

Pc	Axiomatization of Propositional Calculus
K	$K(\varphi \to \psi) \to (K\varphi \to K\psi)$
T	$K\varphi \to \varphi$
4	$K\varphi \to KK\varphi$
5	$\neg K\varphi \to K\neg K\varphi$
MP	$\dfrac{\varphi \qquad \varphi \to \psi}{\psi}$
Nec	$\dfrac{\varphi}{K\psi}$

Theorem

S5 *is sound and strongly complete with respect to the class of Kripke frames with equivalence relations.*

Multi-agent Epistemic Logic

The Language: $\varphi \;:=\; p \mid \neg\varphi \mid \varphi \wedge \psi \mid K\varphi$

Kripke Models: $\mathcal{M} = \langle W, R, V \rangle$ and $w \in W$

Truth: $\mathcal{M}, w \models \varphi$ is defined as follows:

- $\mathcal{M}, w \models p$ iff $w \in V(p)$ (with $p \in \text{At}$)
- $\mathcal{M}, w \models \neg\varphi$ if $\mathcal{M}, w \not\models \varphi$
- $\mathcal{M}, w \models \varphi \wedge \psi$ if $\mathcal{M}, w \models \varphi$ and $\mathcal{M}, w \models \psi$
- $\mathcal{M}, w \models K\varphi$ if for each $v \in W$, if wRv, then $\mathcal{M}, v \models \varphi$

Multi-agent Epistemic Logic

The Language: $\varphi \ := \ p \mid \neg\varphi \mid \varphi \wedge \psi \mid K_i\varphi$ with $i \in \mathcal{A}$

Kripke Models: $\mathcal{M} = \langle W, \{R_i\}_{i \in \mathcal{A}}, V \rangle$ and $w \in W$

Truth: $\mathcal{M}, w \models \varphi$ is defined as follows:

- $\mathcal{M}, w \models p$ iff $w \in V(p)$ (with $p \in \mathrm{At}$)
- $\mathcal{M}, w \models \neg\varphi$ if $\mathcal{M}, w \not\models \varphi$
- $\mathcal{M}, w \models \varphi \wedge \psi$ if $\mathcal{M}, w \models \varphi$ and $\mathcal{M}, w \models \psi$
- $\mathcal{M}, w \models K_i\varphi$ if for each $v \in W$, if wR_iv, then $\mathcal{M}, v \models \varphi$

Multi-agent Epistemic Logic

- $K_A K_B \varphi$: "Ann knows that Bob knows φ"
- $K_A(K_B \varphi \vee K_B \neg\varphi)$: "Ann knows that Bob knows whether φ
- $\neg K_B K_A K_B(\varphi)$: "Bob does not know that Ann knows that Bob knows that φ"

Next:

- More about the logic (complexity results, model checking, Tableaux systems, etc.)
- Group notions
- Additional informational attitudes (belief, certainty, justification, only knowing, awareness, etc.)
- Other semantics