Logic and Artificial Intelligence

Lecture 12

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

October 10, 2011

Dynamic Epistemic Logic: Literature

A. Baltag and L. Moss. *Logics for Epistemic Programs*. 2004.

W. van der Hoek, H. van Ditmarsch and B. Kooi. *Dynamic Epistemic Logic*. 2007.

Dynamic Epistemic Logic

Dynamic Epistemic Logic

$$\mathcal{M} \otimes E$$

Dynamic Epistemic Logic

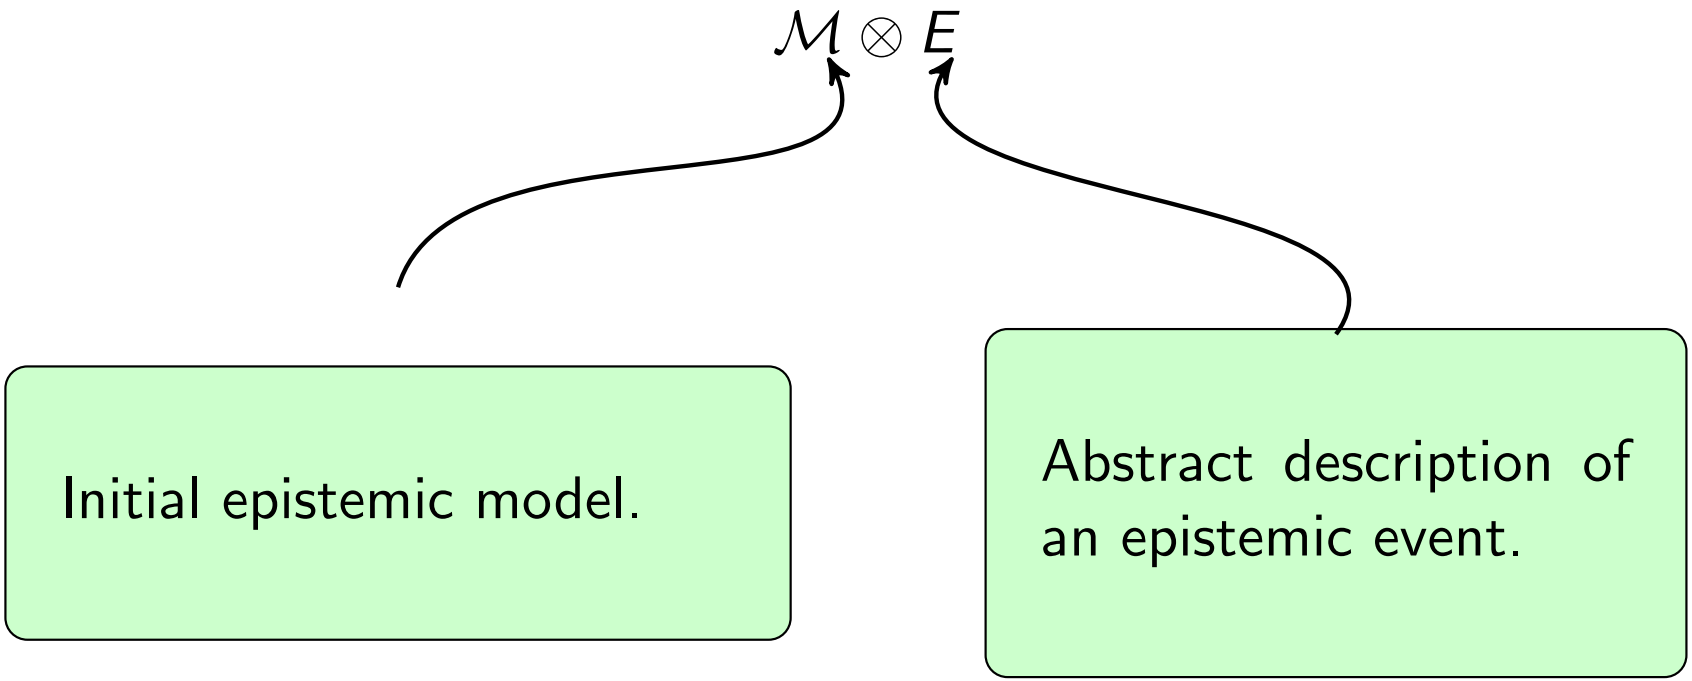

Abstract Description of the Event

Ann looks at the card while Bob is looking away.

Abstract Description of the Event

Ann looks at the card while Bob is looking away.

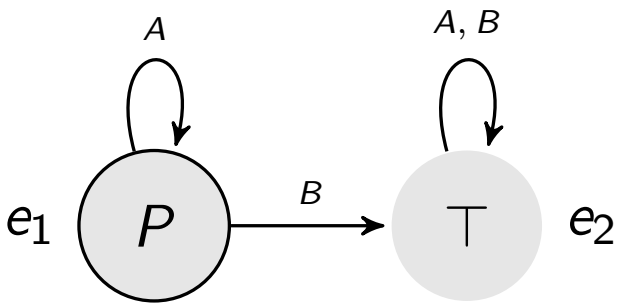

Product Update

Product Update

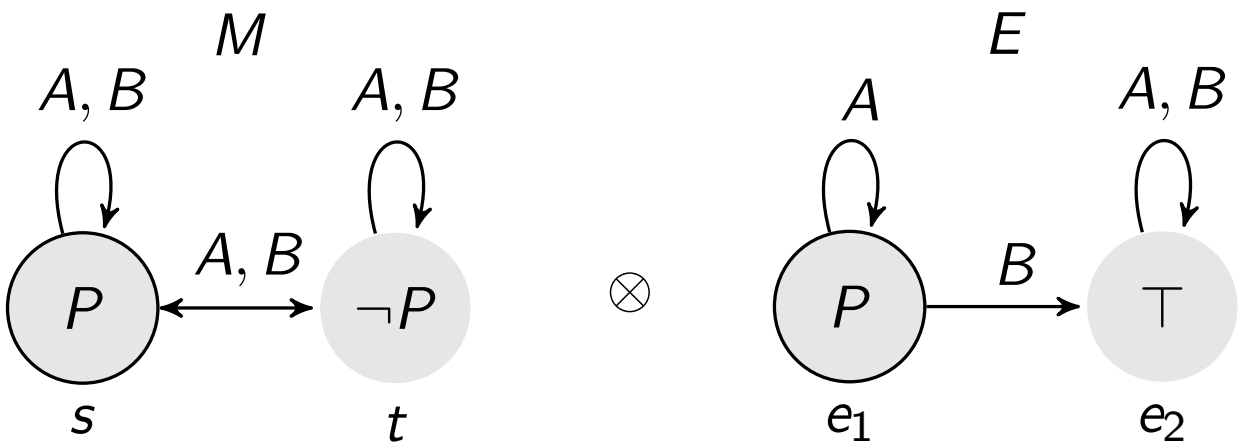

Product Update

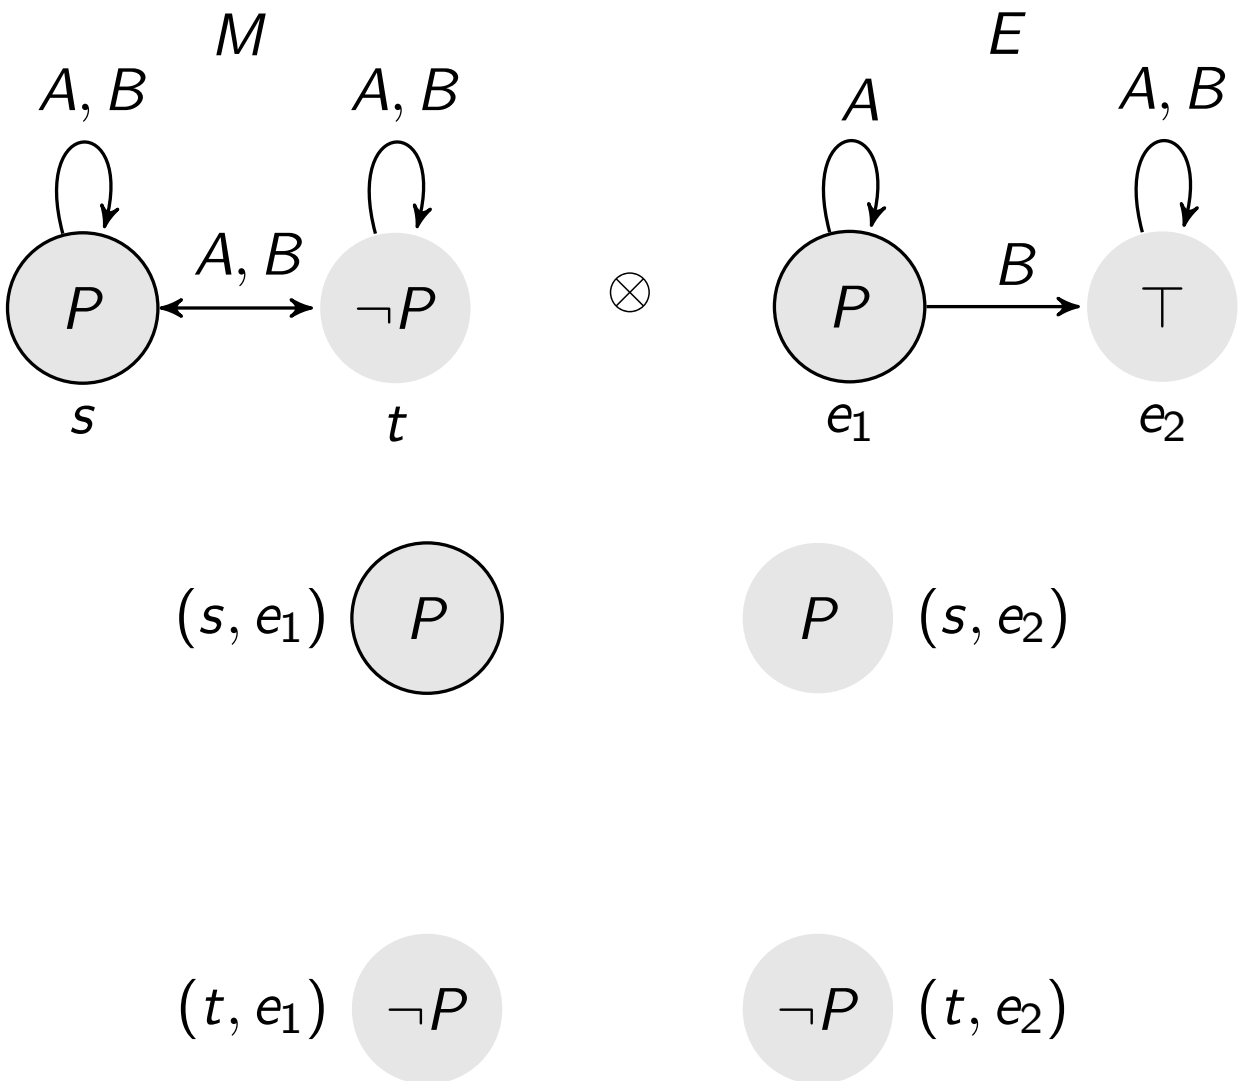

Product Update

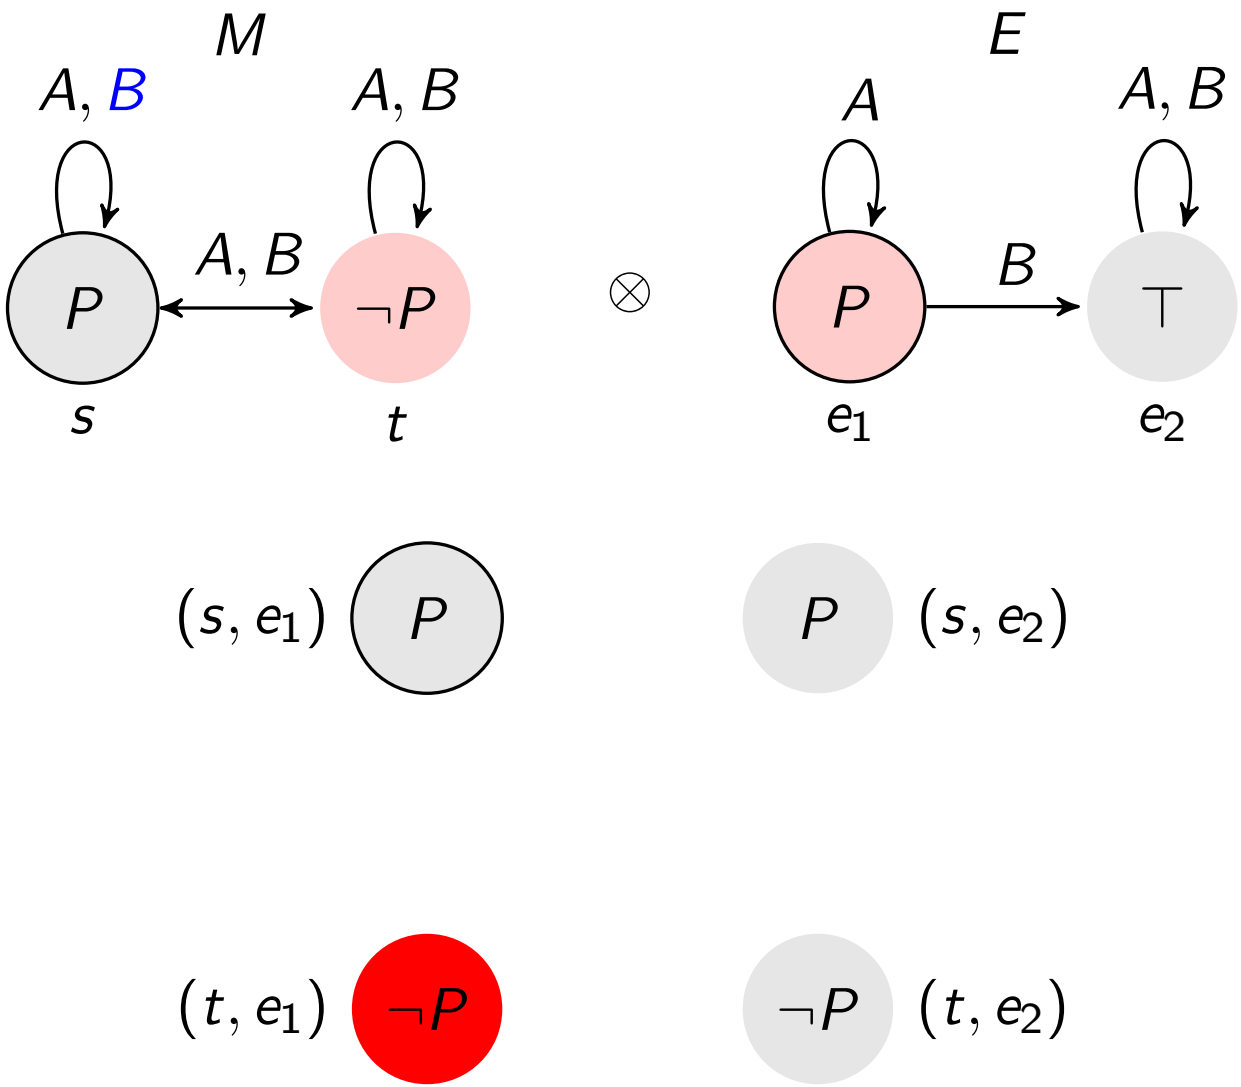

Product Update

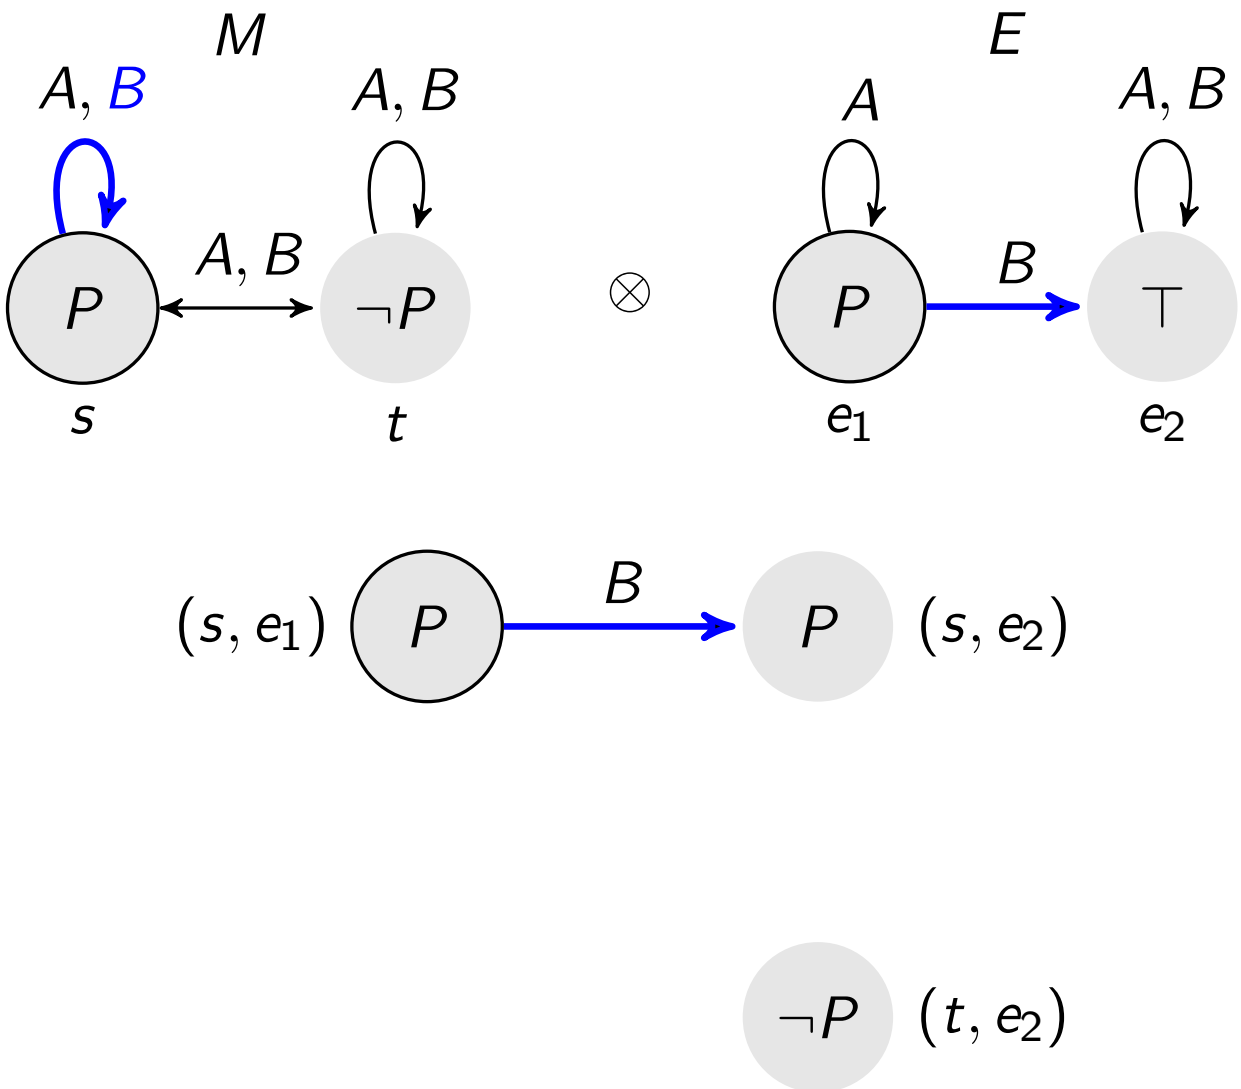

Product Update

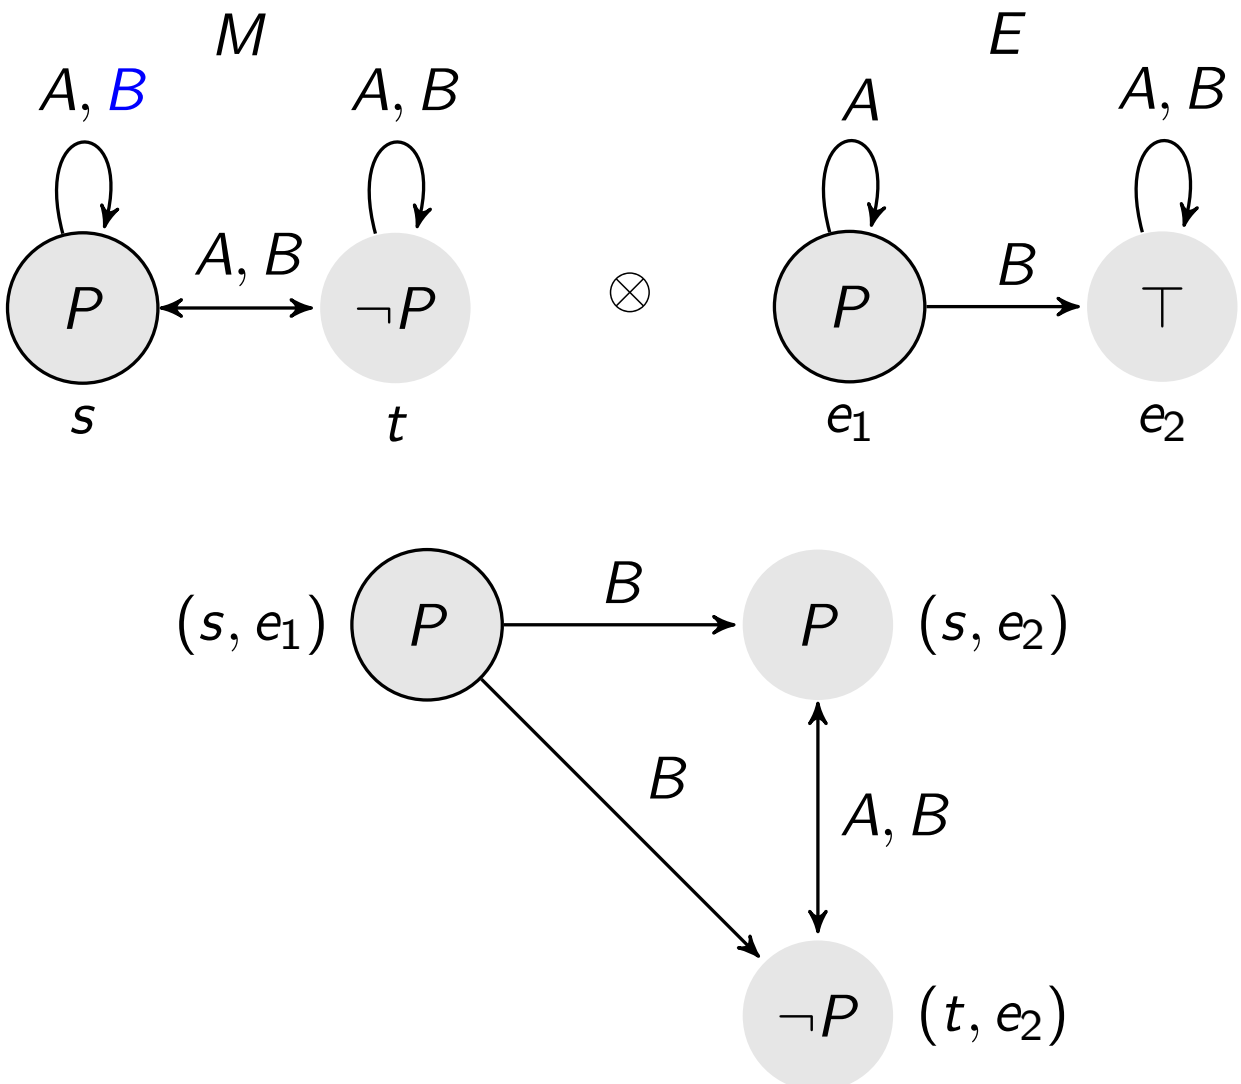

Product Update Details

Let $\mathbb{M} = \langle W, R, V \rangle$ be a Kripke model.

An event model is a tuple $\mathbb{A} = \langle A, S, Pre \rangle$, where $S \subseteq A \times A$ and $Pre : \mathcal{L} \to \wp(A)$.

Product Update Details

Let $\mathbb{M} = \langle W, R, V \rangle$ be a Kripke model.

An event model is a tuple $\mathbb{A} = \langle A, S, Pre \rangle$, where $S \subseteq A \times A$ and $Pre : \mathcal{L} \to \wp(A)$.

The update model $\mathbb{M} \otimes \mathbb{A} = \langle W', R', V' \rangle$ where

Product Update Details

Let $\mathbb{M} = \langle W, R, V \rangle$ be a Kripke model.

An event model is a tuple $\mathbb{A} = \langle A, S, Pre \rangle$, where $S \subseteq A \times A$ and $Pre : \mathcal{L} \to \wp(A)$.

The update model $\mathbb{M} \otimes \mathbb{A} = \langle W', R', V' \rangle$ where

- $W' = \{(w, a) \mid w \models Pre(a)\}$

Product Update Details

Let $\mathbb{M} = \langle W, R, V \rangle$ be a Kripke model.

An event model is a tuple $\mathbb{A} = \langle A, S, Pre \rangle$, where $S \subseteq A \times A$ and $Pre : \mathcal{L} \to \wp(A)$.

The update model $\mathbb{M} \otimes \mathbb{A} = \langle W', R', V' \rangle$ where

- $W' = \{(w, a) \mid w \models Pre(a)\}$
- $(w, a)R'(w', a')$ iff wRw' and aSa'

Product Update Details

Let $\mathbb{M} = \langle W, R, V \rangle$ be a Kripke model.

An event model is a tuple $\mathbb{A} = \langle A, S, Pre \rangle$, where $S \subseteq A \times A$ and $Pre : \mathcal{L} \to \wp(A)$.

The update model $\mathbb{M} \otimes \mathbb{A} = \langle W', R', V' \rangle$ where

- $W' = \{(w, a) \mid w \models Pre(a)\}$
- $(w, a)R'(w', a')$ iff wRw' and aSa'
- $(w, a) \in V'(p)$ iff $w \in V(p)$

Product Update Details

Let $\mathbb{M} = \langle W, R, V \rangle$ be a Kripke model.

An event model is a tuple $\mathbb{A} = \langle A, S, Pre \rangle$, where $S \subseteq A \times A$ and $Pre : \mathcal{L} \to \wp(A)$.

The update model $\mathbb{M} \otimes \mathbb{A} = \langle W', R', V' \rangle$ where

- $W' = \{(w, a) \mid w \models Pre(a)\}$
- $(w, a)R'(w', a')$ iff wRw' and aSa'
- $(w, a) \in V'(p)$ iff $w \in V(p)$

$\mathcal{M}, w \models [\mathrm{A}, a]\varphi$ iff $\mathcal{M}, w \models Pre(a)$ implies $\mathcal{M} \otimes A, (w, a) \models \varphi$.

Example

Ann would like Bob to attend her talk; however, she only wants Bob to attend if he is interested in the subject of her talk, not because he is just being polite.

There is a very simple procedure to solve Ann's problem: *have a (trusted) friend tell Bob the time and subject of her talk.*

Is this procedure correct?

Example

Ann would like Bob to attend her talk; however, she only wants Bob to attend if he is interested in the subject of her talk, not because he is just being polite.

There is a very simple procedure to solve Ann's problem: *have a (trusted) friend tell Bob the time and subject of her talk.*

Is this procedure correct? Yes, if

1. Ann knows about the talk.
2. Bob knows about the talk.
3. Ann knows that Bob knows about the talk.
4. Bob *does not* know that Ann knows that he knows about the talk.
5. *And nothing else.*

Dynamic Epistemic Logic

Recall the Ann and Bob example: Charles tells Bob that the talk is at 2PM.

Dynamic Epistemic Logic

Recall the Ann and Bob example: Charles tells Bob that the talk is at 2PM.

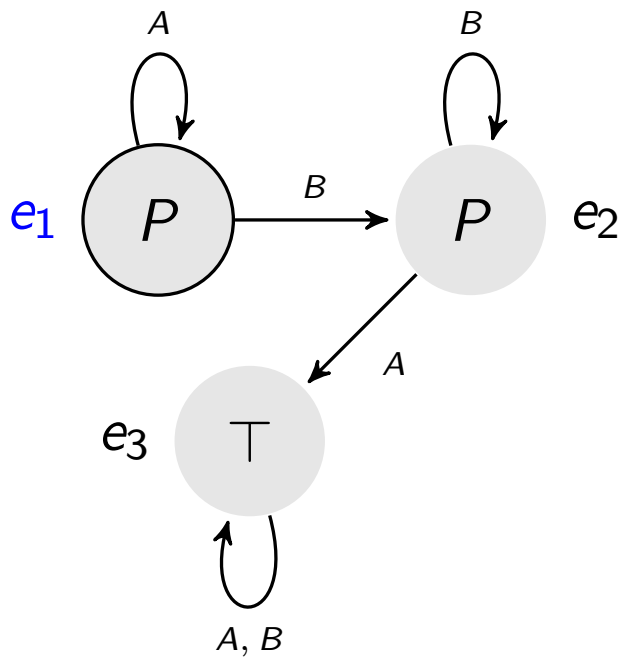

Dynamic Epistemic Logic

Recall the Ann and Bob example: Charles tells Bob that the talk is at 2PM.

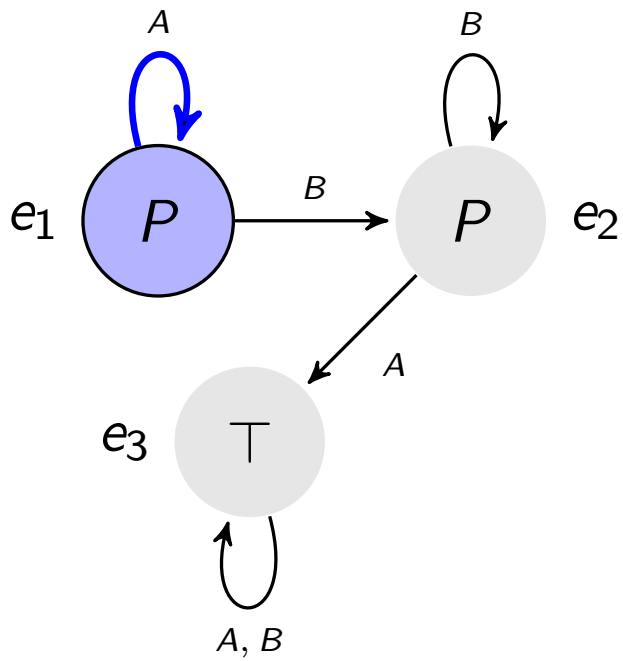

Ann knows which event took place.

Dynamic Epistemic Logic

Recall the Ann and Bob example: Charles tells Bob that the talk is at 2PM.

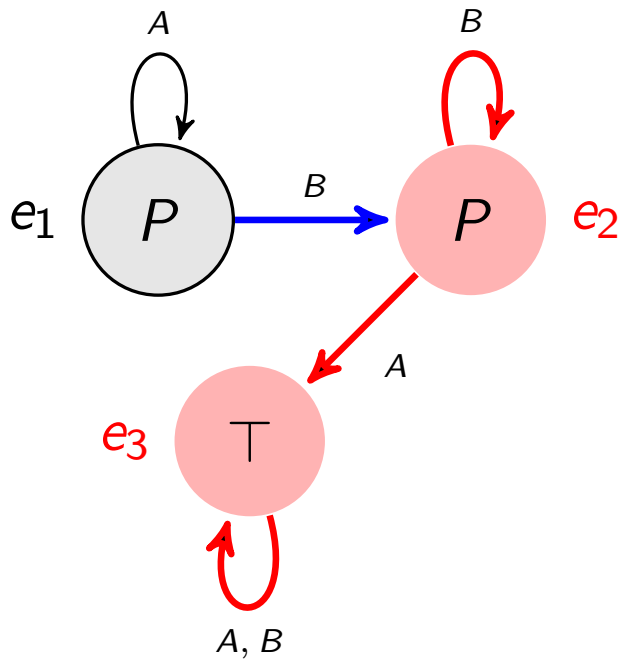

Bob thinks a different event took place.

Dynamic Epistemic Logic

Recall the Ann and Bob example: Charles tells Bob that the talk is at 2PM.

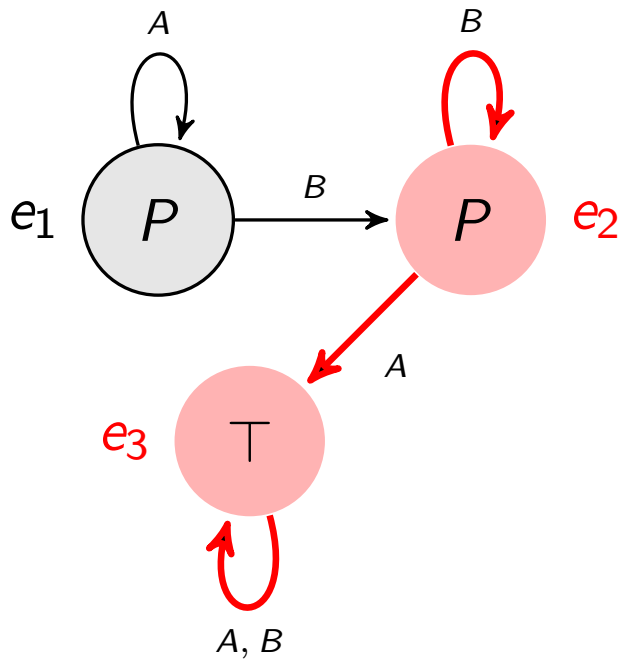

That is, Bob learns the time of the talk, but Ann learns nothing.

Dynamic Epistemic Logic

Dynamic Epistemic Logic

$$(\mathcal{M} \otimes E_1) \otimes E_2$$

Dynamic Epistemic Logic

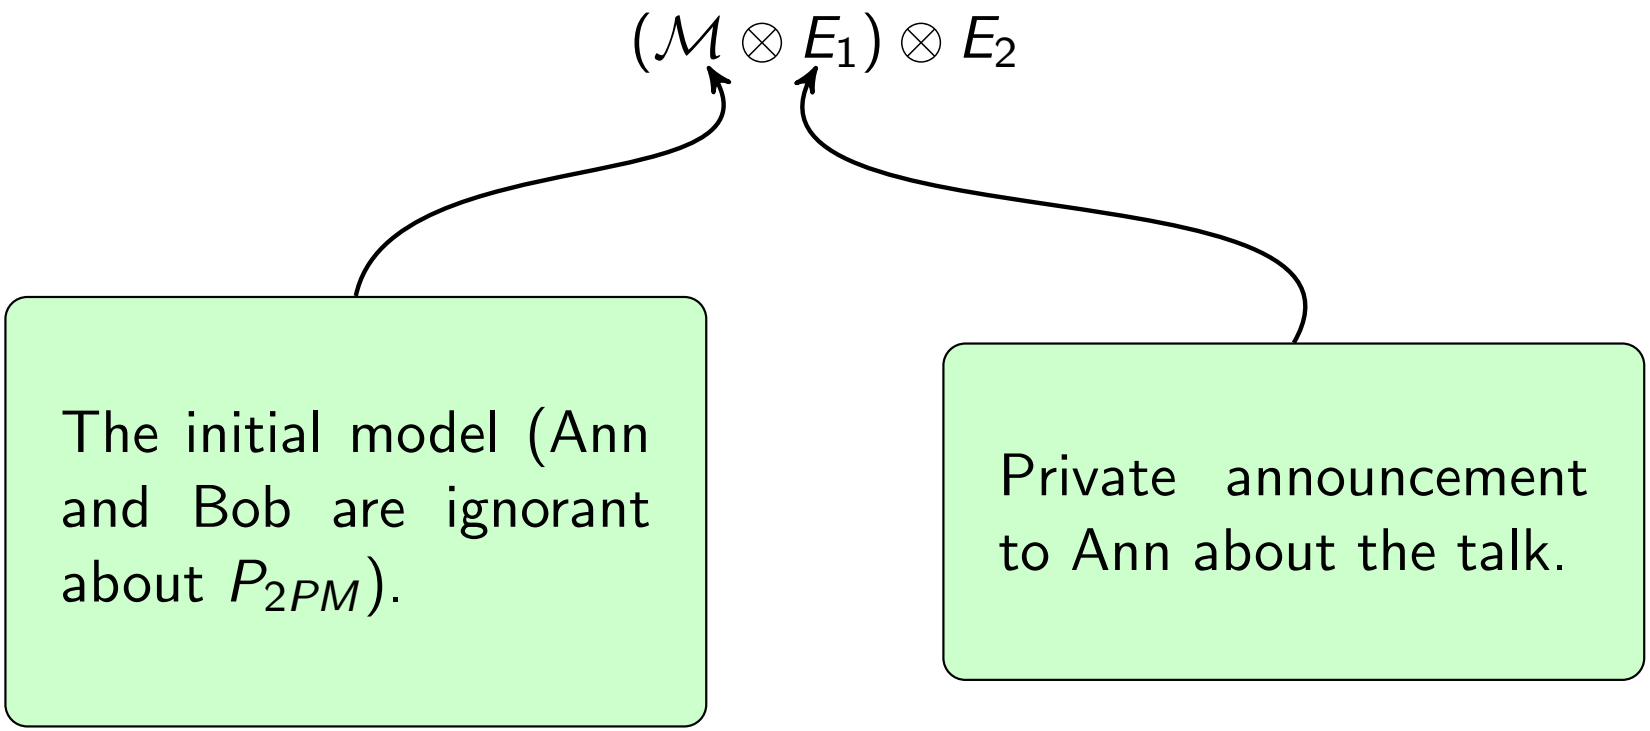

Product Update

Product Update

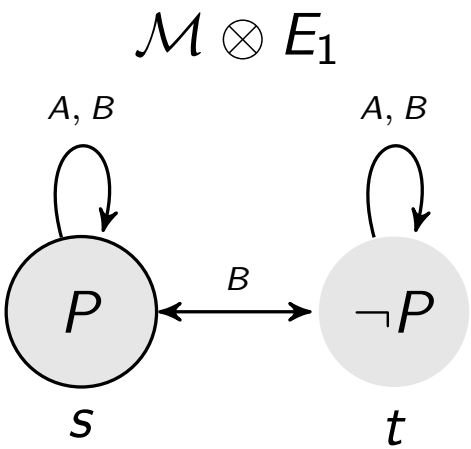

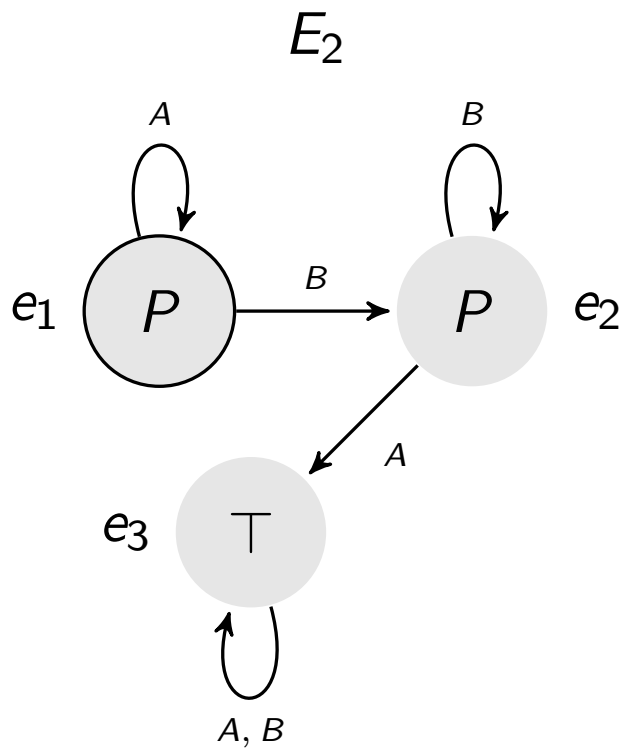

Product Update

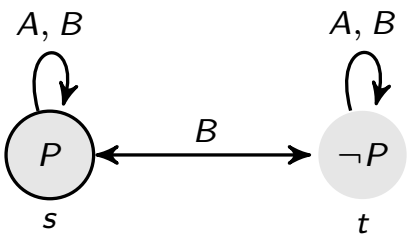

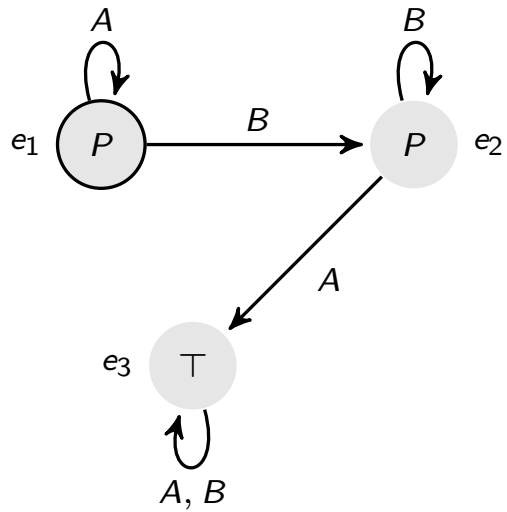

(s, e_1) P P (s, e_2)

(s, e_3) P $\neg P$ (t, e_3)

Product Update

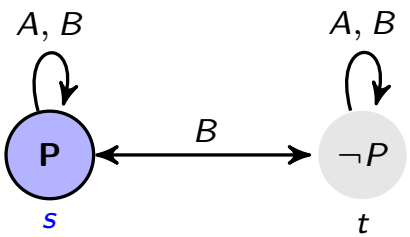

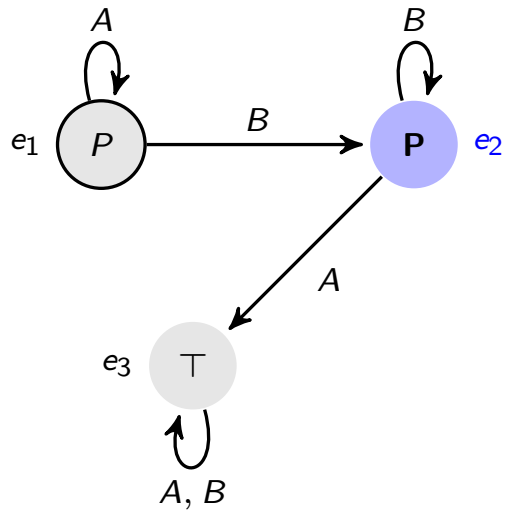

(s, e_1) P

P (s, e_2)

(s, e_3) P

$\neg P$ (t, e_3)

Product Update

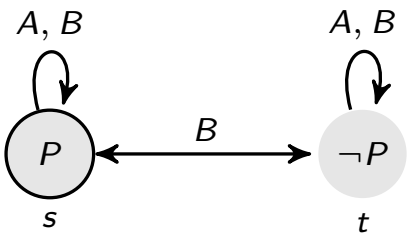

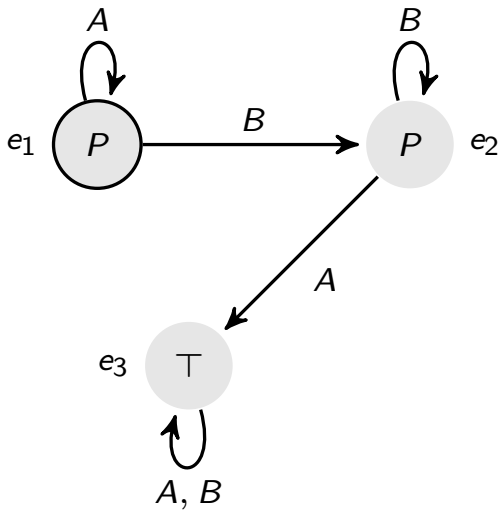

$(s, e_1) \models \neg K_B K_A K_B P$ (s, e_1) P P (s, e_2)

(s, e_3) P $\neg P$ (t, e_3)

Product Update

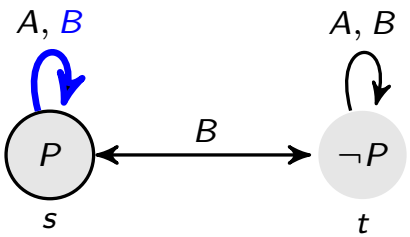

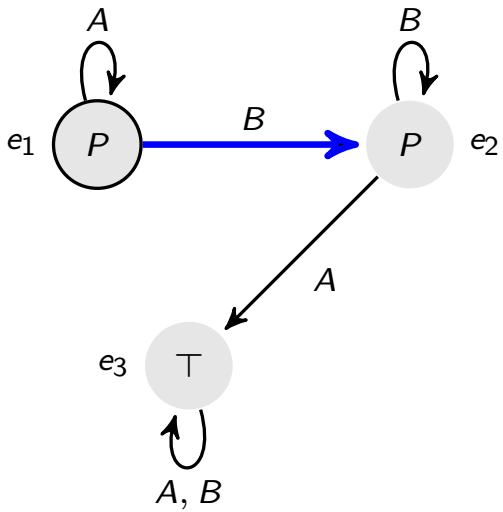

$(s, e_1) \models \neg K_B K_A K_B P$ (s, e_1) P B P (s, e_2)

(s, e_3) P $\neg P$ (t, e_3)

Product Update

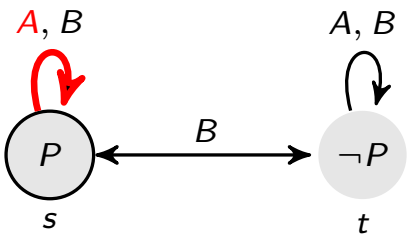

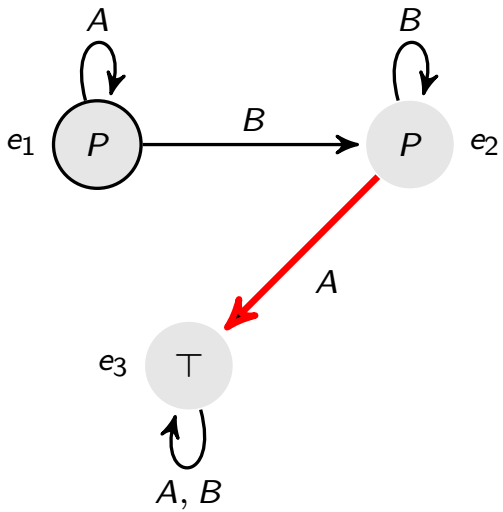

$(s, e_1) \models \neg K_B K_A K_B P$

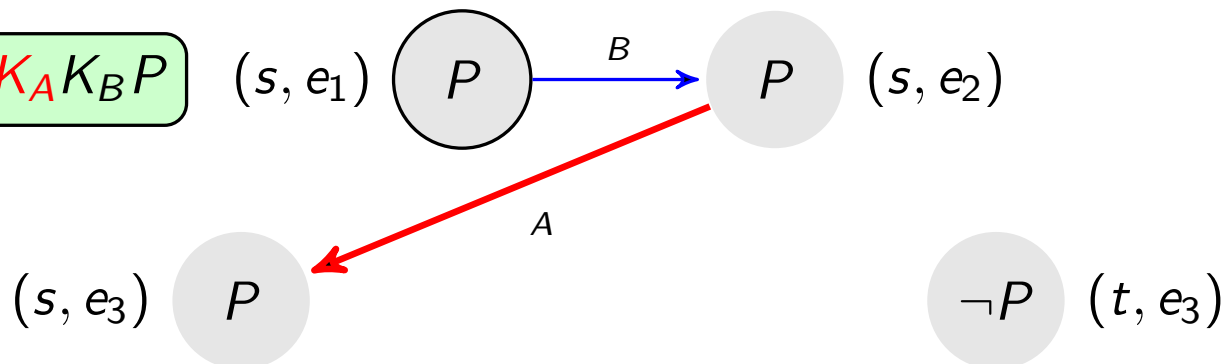

Product Update

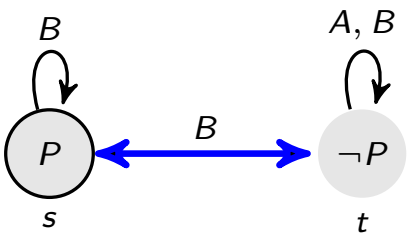

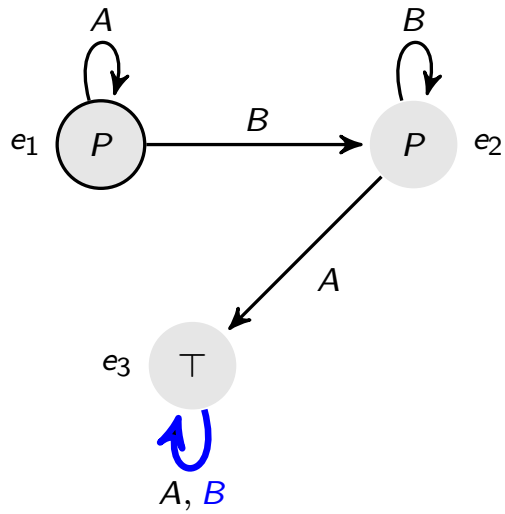

$(s, e_1) \models \neg K_B K_A K_B P$

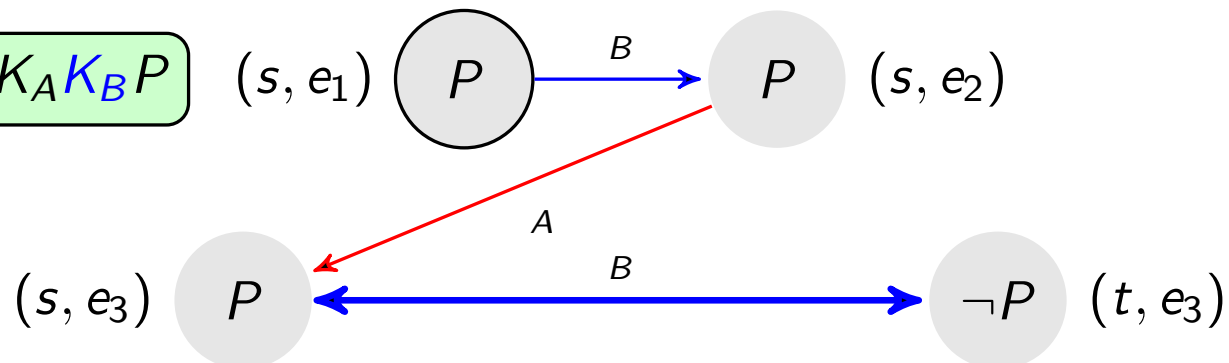

1. The agents' *observational* powers.

2. The *type* of change triggered by the event.

3. The underlying *protocol* specifying which events (observations, messages, actions) are available (or permitted) at any given moment.

1. The agents' *observational* powers.

2. The *type* of change triggered by the event.

3. The underlying *protocol* specifying which events (observations, messages, actions) are available (or permitted) at any given moment.

A recurring issue in any formal model representing agents' (changing) informational attitudes is how to account for the fact that the agents are limited in their access to the available inference steps, possible observations and available messages.

A recurring issue in any formal model representing agents' (changing) informational attitudes is how to account for the fact that the agents are limited in their access to the available inference steps, possible observations and available messages.

1. This may be because the agents are not *logically omniscient* and so do not have unlimited reasoning ability.

A recurring issue in any formal model representing agents' (changing) informational attitudes is how to account for the fact that the agents are limited in their access to the available inference steps, possible observations and available messages.

1. This may be because the agents are not *logically omniscient* and so do not have unlimited reasoning ability.

2. But it can also be because the agents are following a predefined **protocol** that *explicitly or implicitly limits* statements available for observation and/or communication.

A recurring issue in any formal model representing agents' (changing) informational attitudes is how to account for the fact that the agents are limited in their access to the available inference steps, possible observations and available messages.

1. This may be because the agents are not *logically omniscient* and so do not have unlimited reasoning ability.

✓ But it can also be because the agents are following a predefined **protocol** that *explicitly or implicitly limits* statements available for observation and/or communication.

Epistemic Temporal Logic

R. Parikh and R. Ramanujam. *A Knowledge Based Semantics of Messages*. *Journal of Logic, Language and Information*, 12: 453 – 467, 1985, 2003.

FHMV. Reasoning about Knowledge. MIT Press, 1995.

The 'Playground'

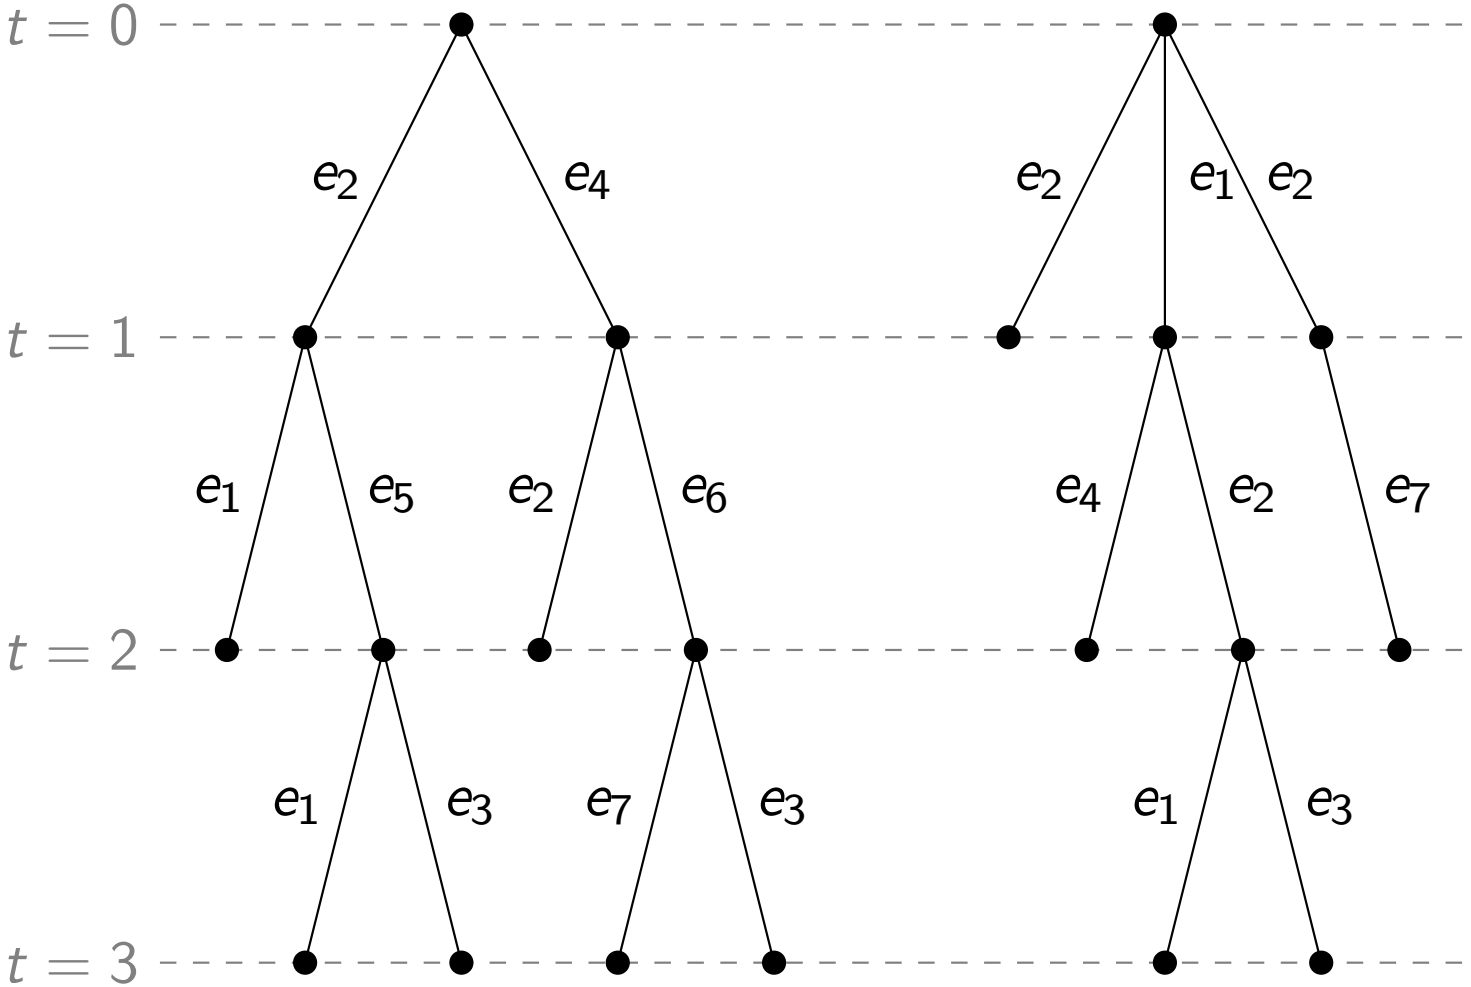

The 'Playground'

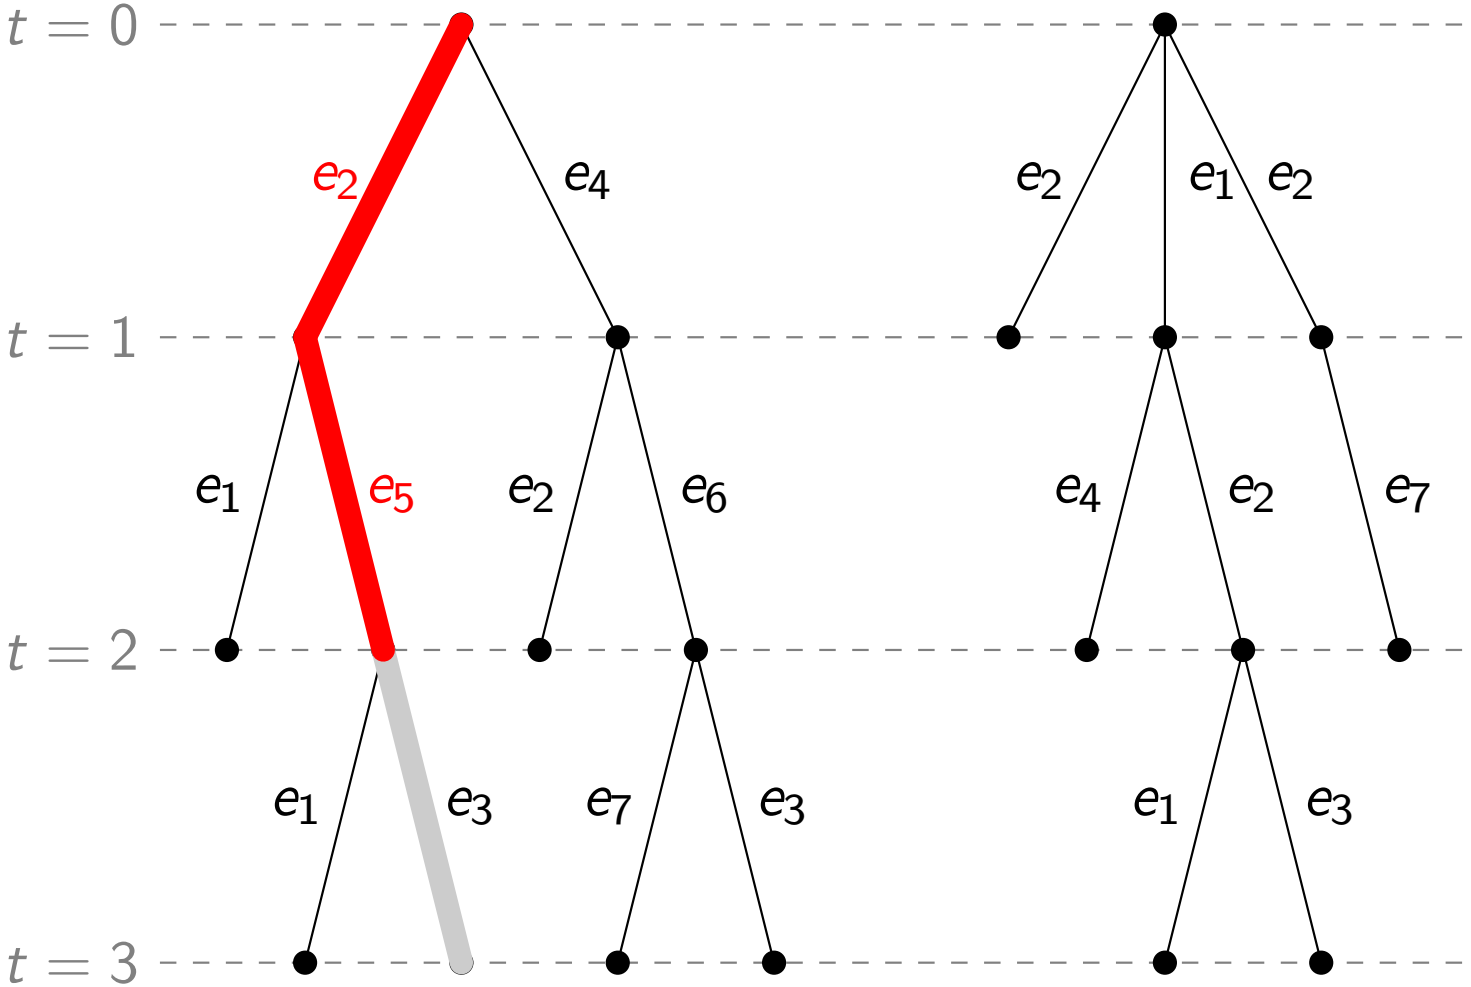

The 'Playground'

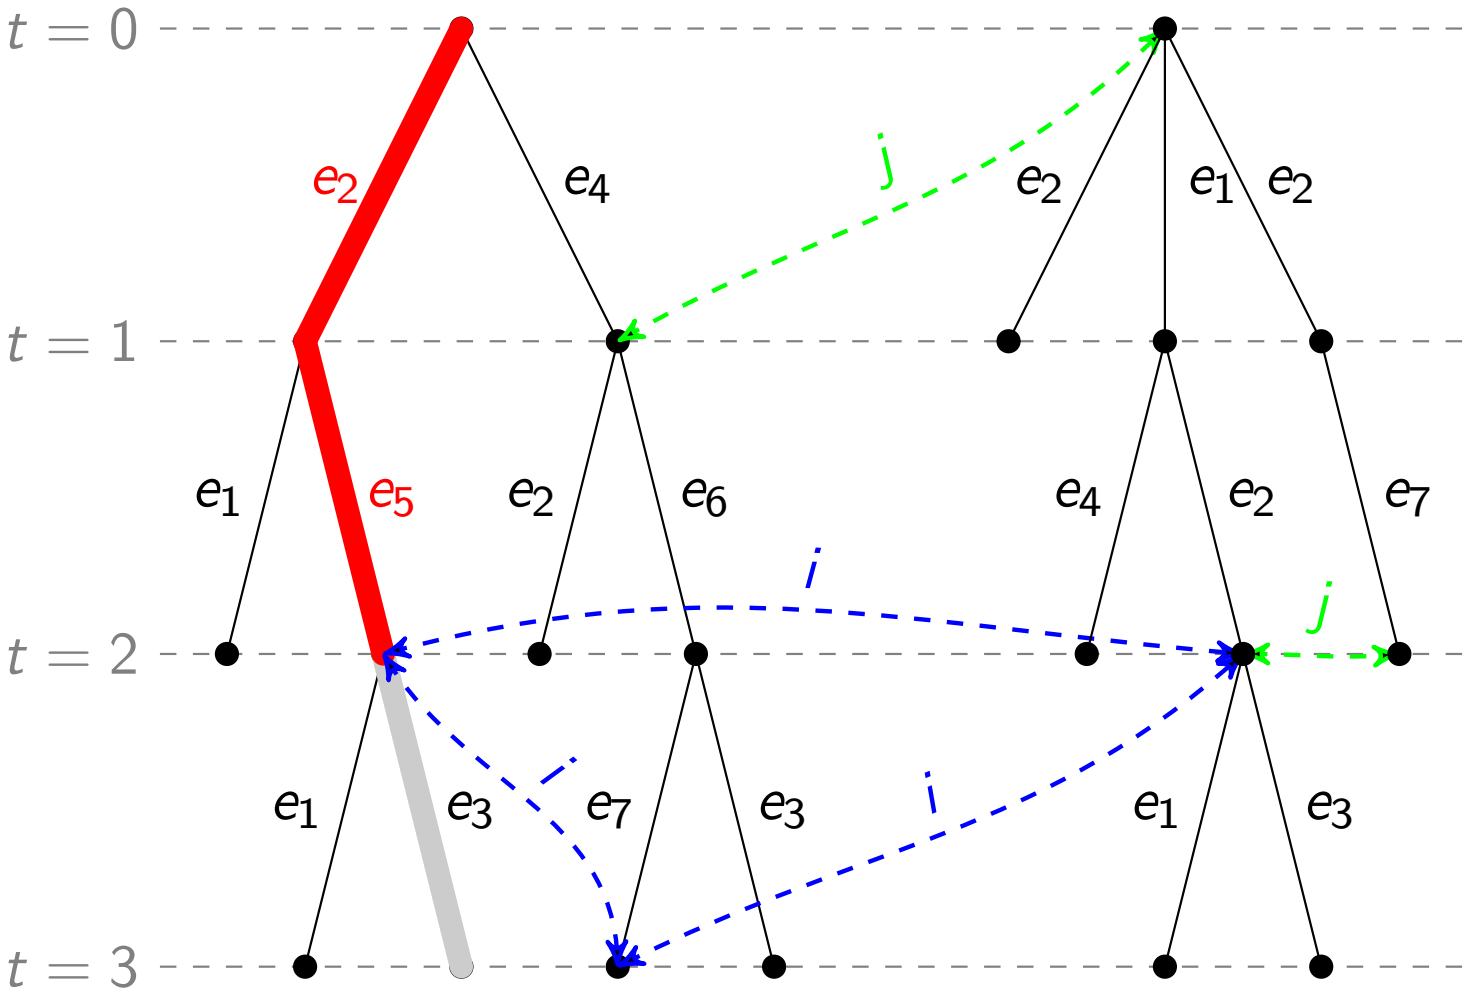

Formal Languages

- $P\varphi$ (φ is true *sometime* in the past),
- $F\varphi$ (φ is true *sometime* in the future),
- $Y\varphi$ (φ is true at *the* previous moment),
- $N\varphi$ (φ is true at *the* next moment),
- $N_e\varphi$ (φ is true after event e)
- $K_i\varphi$ (agent i knows φ) and
- $C_B\varphi$ (the group $B \subseteq \mathcal{A}$ commonly knows φ).

History-based Models

An ETL **model** is a structure $\langle \mathcal{H}, \{\sim_i\}_{i\in\mathcal{A}}, V\rangle$ where $\langle \mathcal{H}, \{\sim_i\}_{i\in\mathcal{A}}\rangle$ is an ETL frame and

$V : \mathrm{At} \to 2^{\mathsf{finite}(\mathcal{H})}$ is a valuation function.

Formulas are interpreted at pairs H, t:

$$H, t \models \varphi$$

Truth in a Model

- $H, t \models P\varphi$ iff there exists $t' \leq t$ such that $H, t' \models \varphi$
- $H, t \models F\varphi$ iff there exists $t' \geq t$ such that $H, t' \models \varphi$
- $H, t \models N\varphi$ iff $H, t+1 \models \varphi$
- $H, t \models Y\varphi$ iff $t > 1$ and $H, t-1 \models \varphi$
- $H, t \models K_i\varphi$ iff for each $H' \in \mathcal{H}$ and $m \geq 0$ if $H_t \sim_i H'_m$ then $H', m \models \varphi$
- $H, t \models C\varphi$ iff for each $H' \in \mathcal{H}$ and $m \geq 0$ if $H_t \sim_* H'_m$ then $H', m \models \varphi$.

where $\sim_*$ is the reflexive transitive closure of the union of the $\sim_i$.

Truth in a Model

- $H, t \models P\varphi$ iff there exists $t' \leq t$ such that $H, t' \models \varphi$
- $H, t \models F\varphi$ iff there exists $t' \geq t$ such that $H, t' \models \varphi$
- $H, t \models N\varphi$ iff $H, t+1 \models \varphi$
- $H, t \models Y\varphi$ iff $t > 1$ and $H, t-1 \models \varphi$
- $H, t \models K_i\varphi$ iff for each $H' \in \mathcal{H}$ and $m \geq 0$ if $H_t \sim_i H'_m$ then $H', m \models \varphi$
- $H, t \models C\varphi$ iff for each $H' \in \mathcal{H}$ and $m \geq 0$ if $H_t \sim_* H'_m$ then $H', m \models \varphi$.

where $\sim_*$ is the reflexive transitive closure of the union of the $\sim_i$.

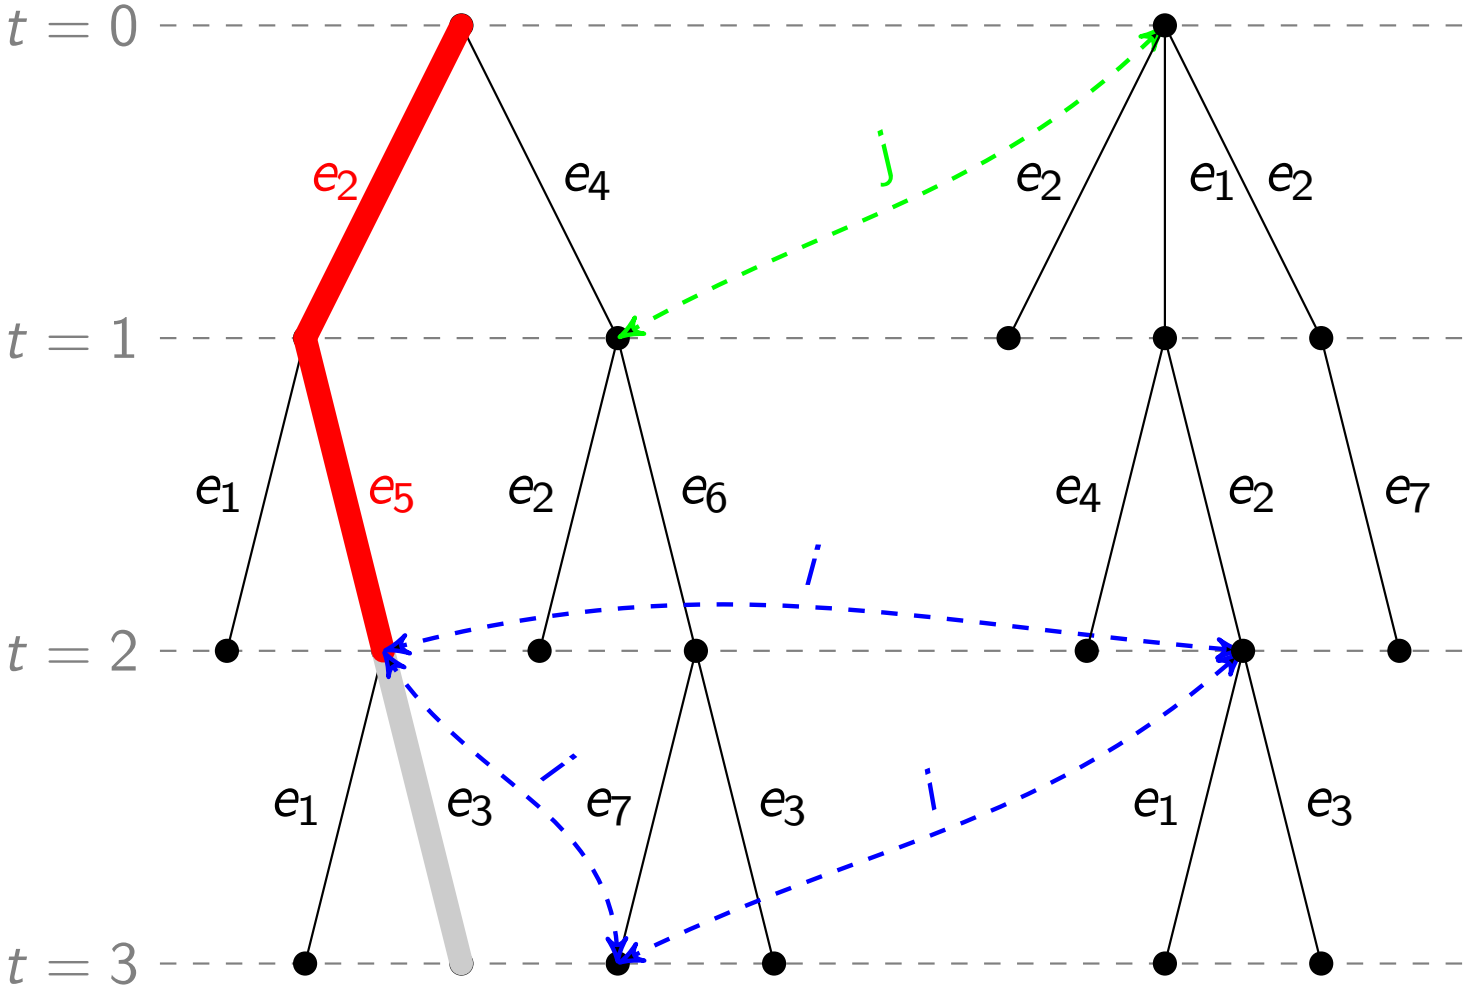
t = 0
t = 1
t = 2
t = 3
e2
e4
j
e2
e1
e2
e1
e5
e2
e6
e4
e2
e7
i
e1
e3
e7
e3
i
e1
e3

Returning to the Example

Ann would like Bob to attend her talk; however, she only wants Bob to attend if he is interested in the subject of her talk, not because he is just being polite.

Returning to the Example

Ann would like Bob to attend her talk; however, she only wants Bob to attend if he is interested in the subject of her talk, not because he is just being polite.

There is a very simple procedure to solve Ann's problem: *have a (trusted) friend tell Bob the time and subject of her talk.*

Returning to the Example

Ann would like Bob to attend her talk; however, she only wants Bob to attend if he is interested in the subject of her talk, not because he is just being polite.

There is a very simple procedure to solve Ann's problem: *have a (trusted) friend tell Bob the time and subject of her talk.*

Is this procedure correct?

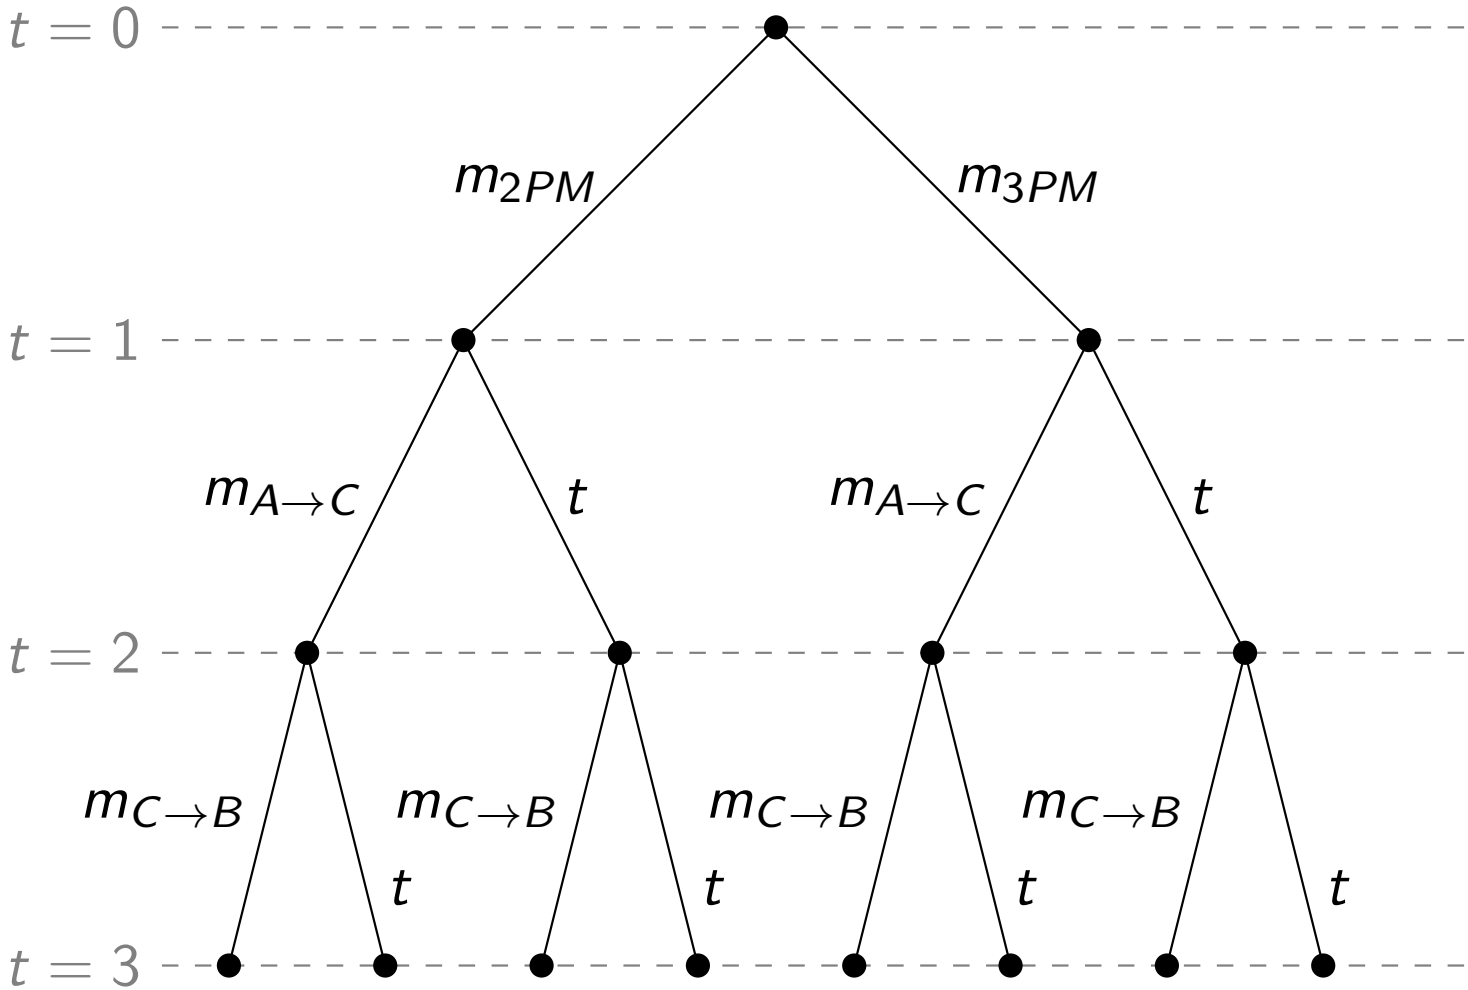
$t = 0$
$t = 1$
$t = 2$
$t = 3$
m_{2PM}
m_{3PM}
$m_{A\to C}$
t
$m_{A\to C}$
t
$m_{C\to B}$
t
$m_{C\to B}$
t
$m_{C\to B}$
t
$m_{C\to B}$
t

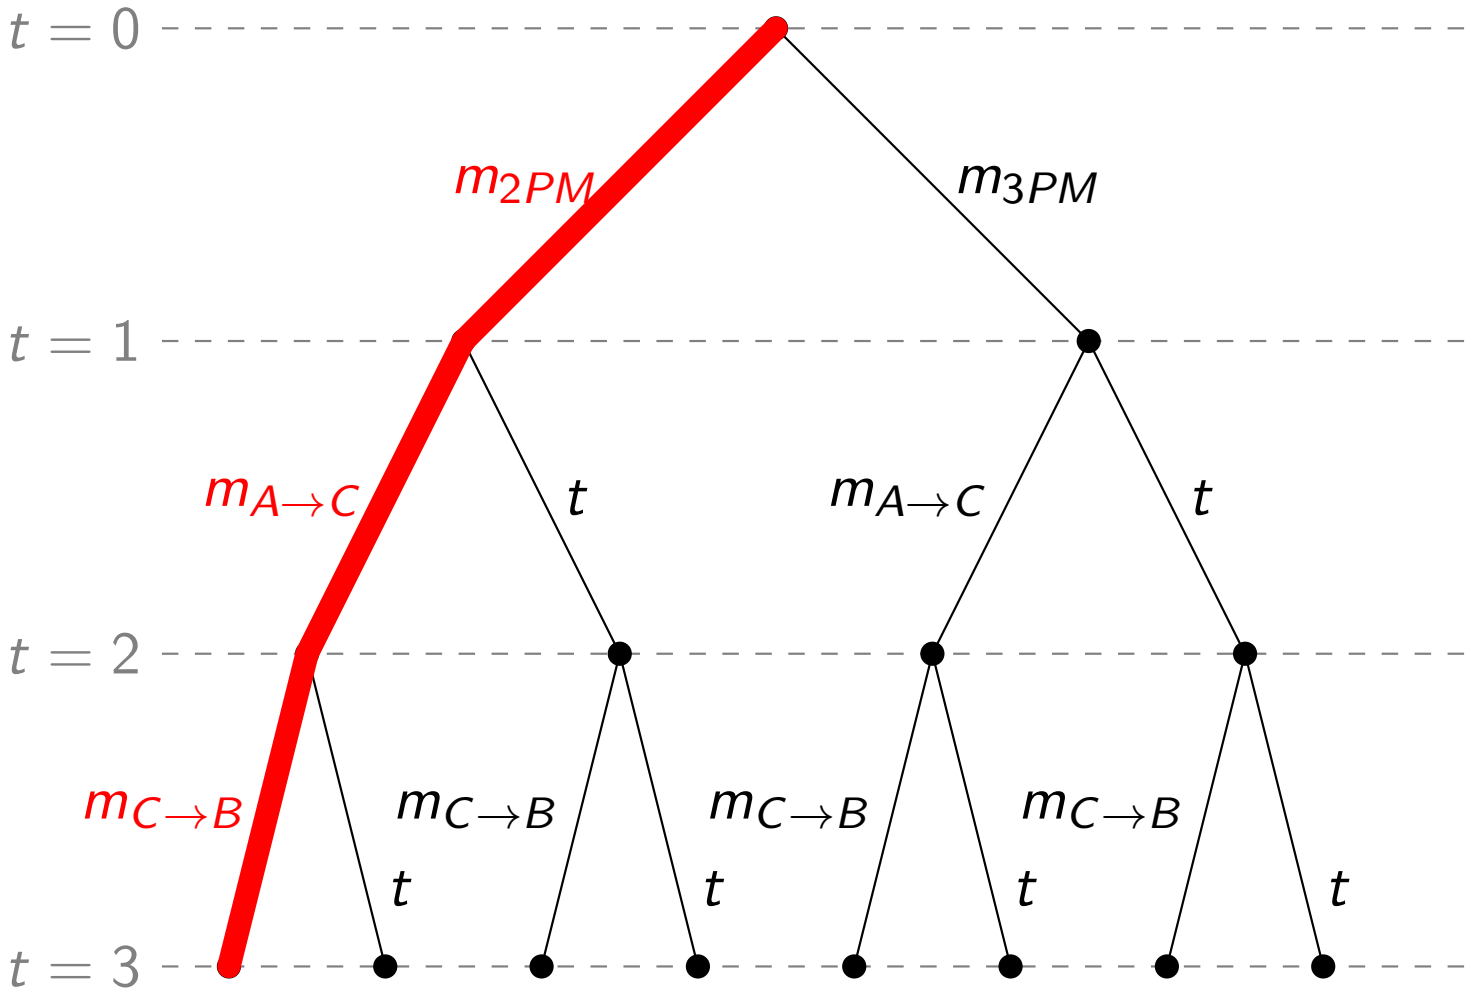

$$H, 3 \models \varphi$$

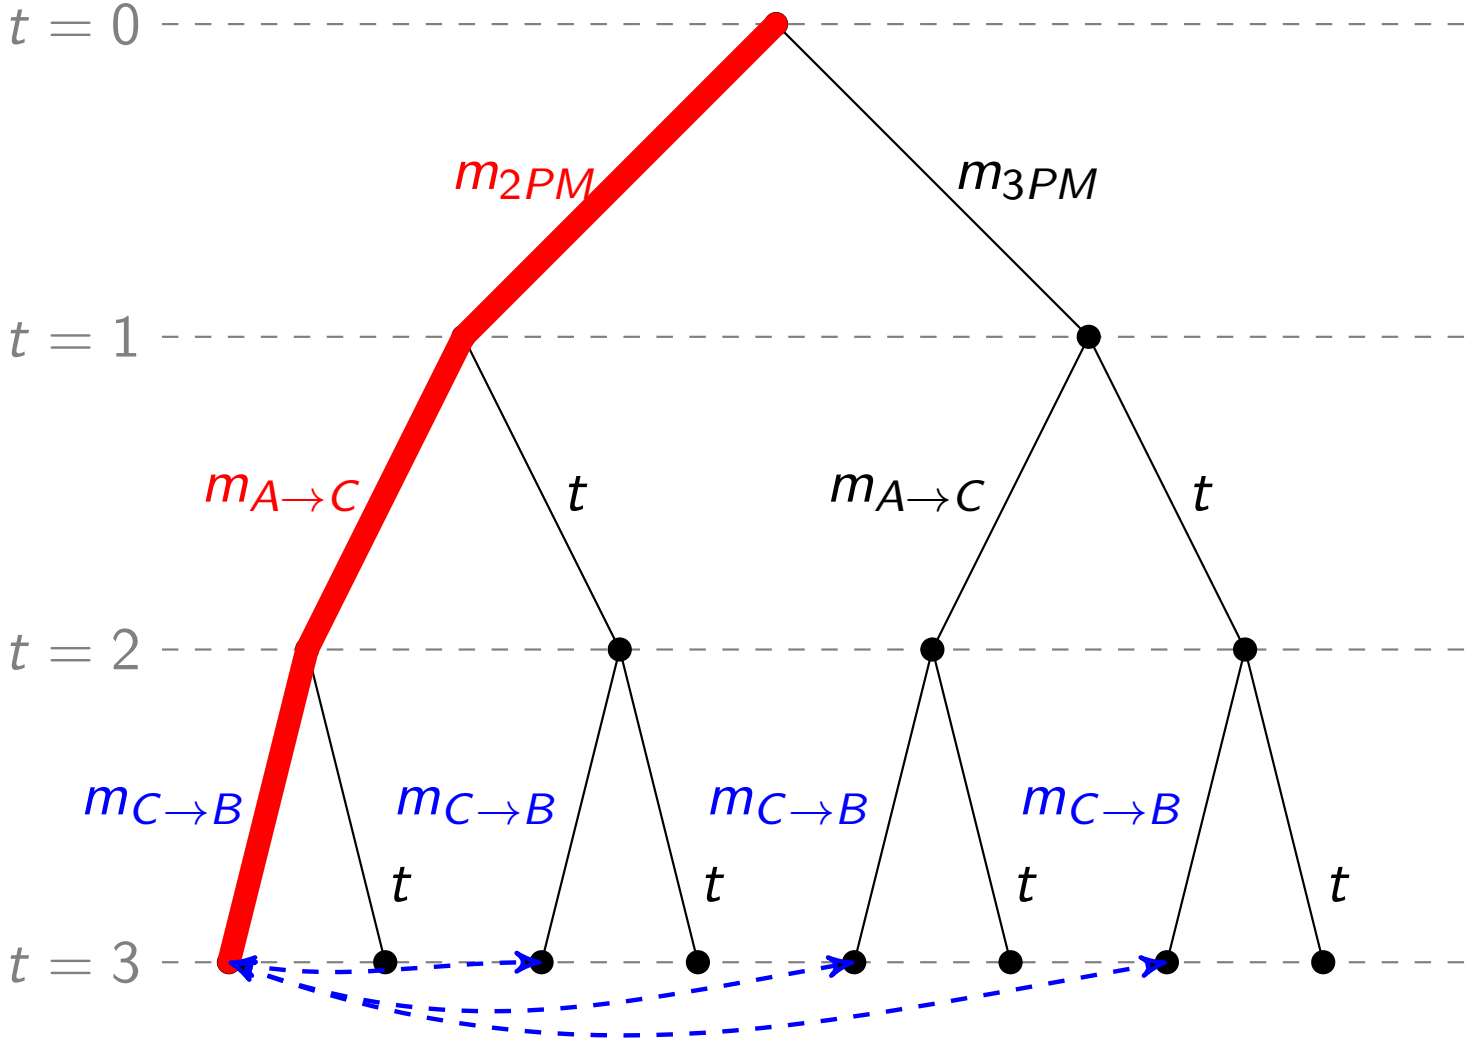

Bob's uncertainty: $H, 3 \models \neg K_B P_{2PM}$

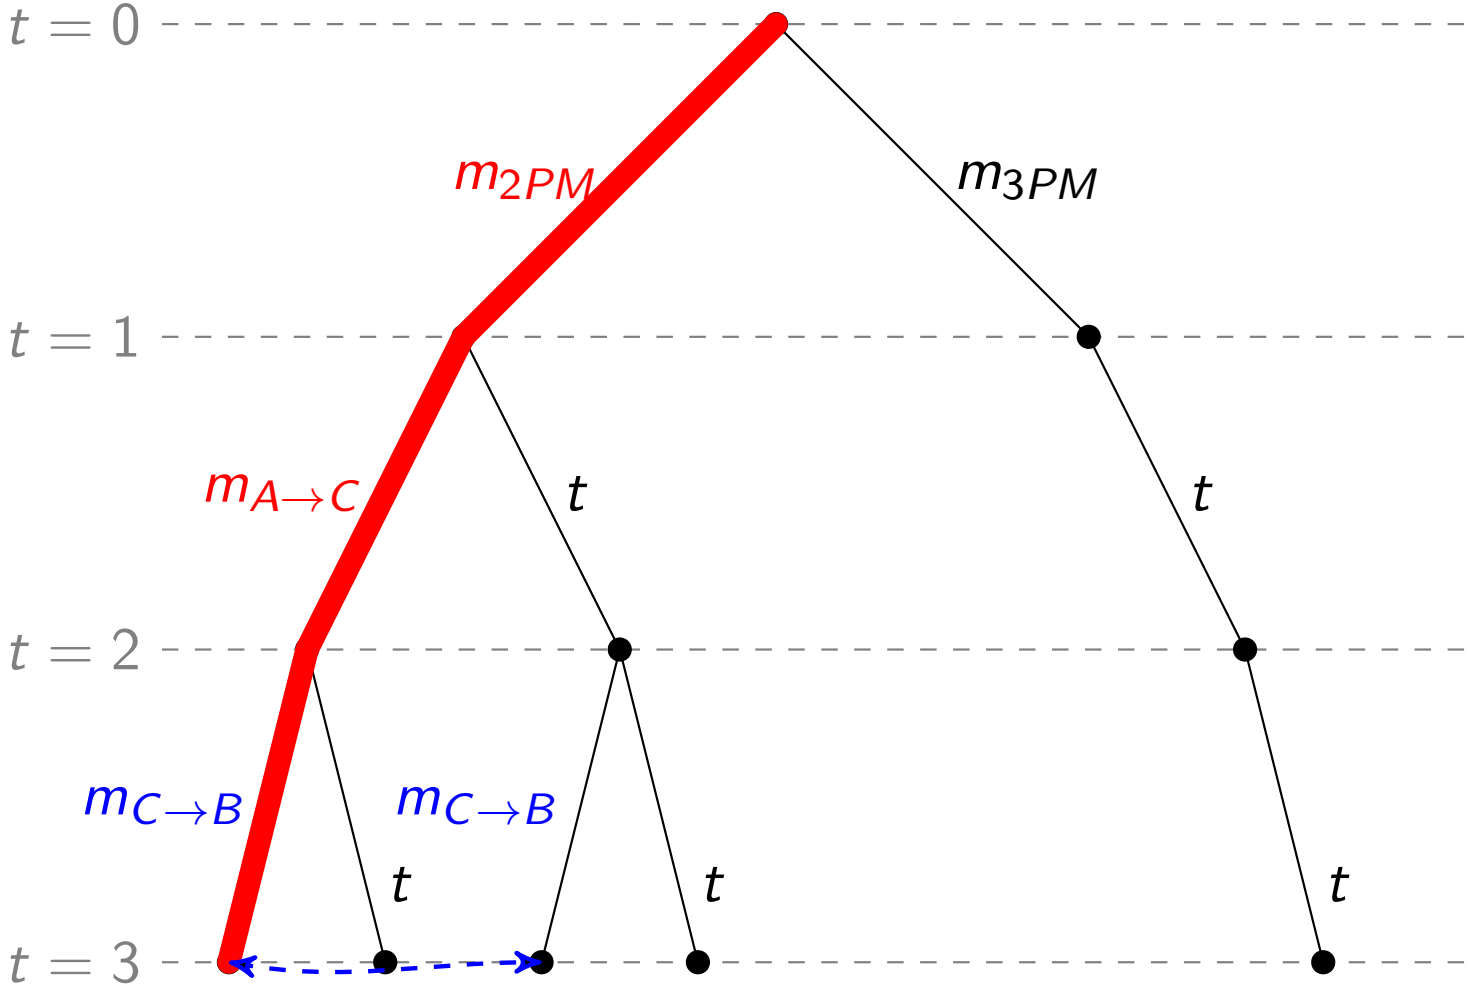

Bob's uncertainty + 'Protocol information': $H, 3 \models K_B P_{2PM}$

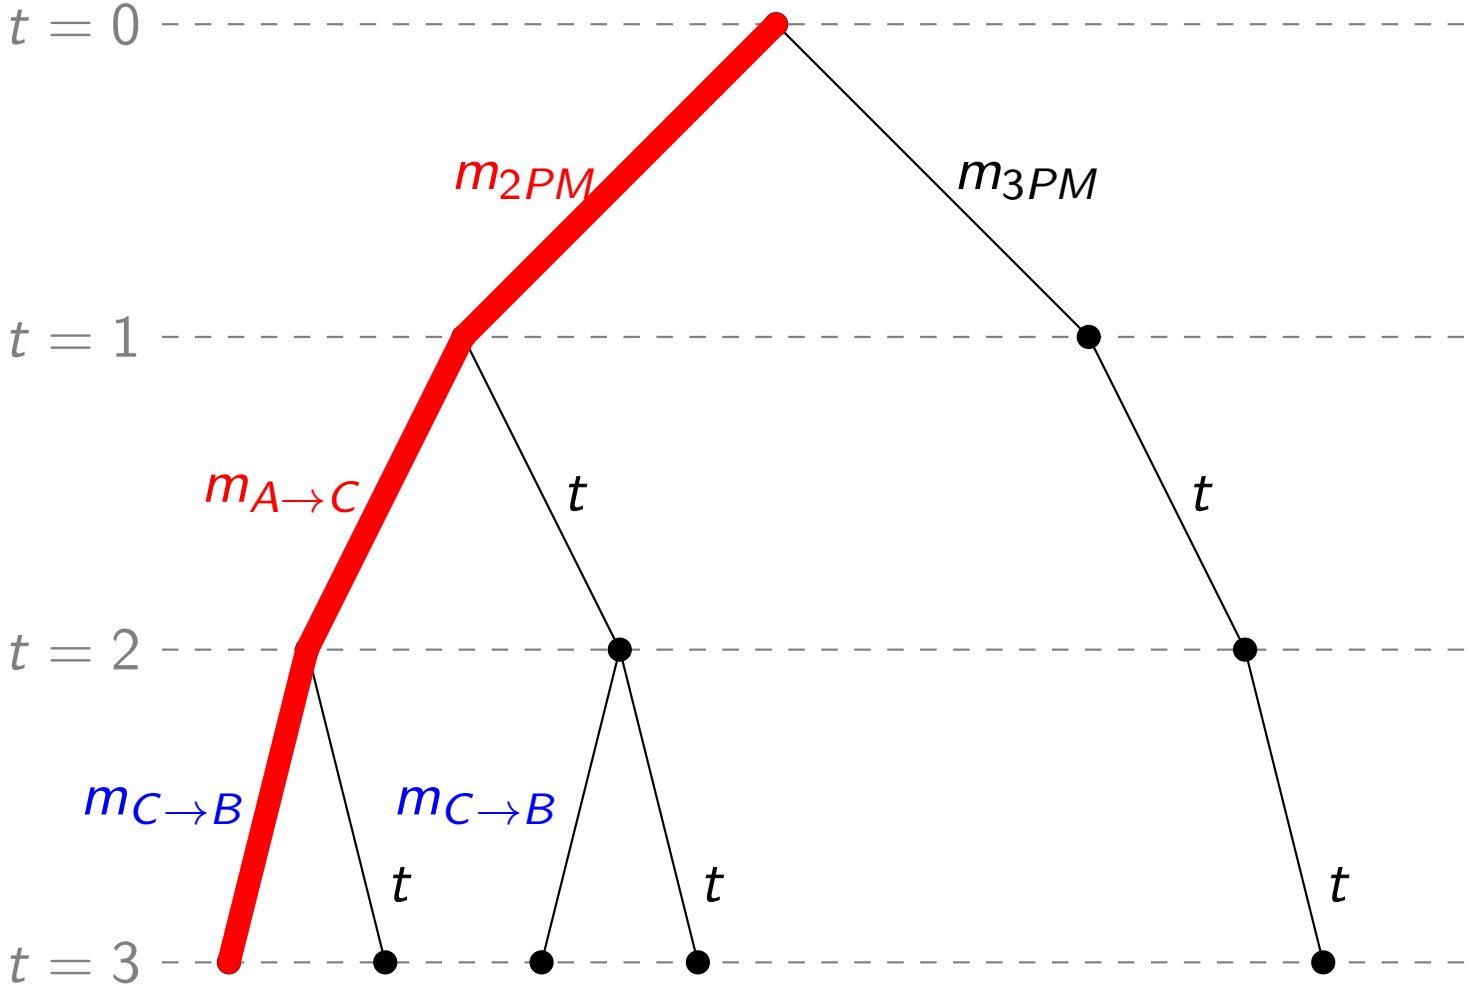

Bob's uncertainty + 'Protocol information':

$$H, 3 \models \neg K_B K_A K_B P_{2PM}$$

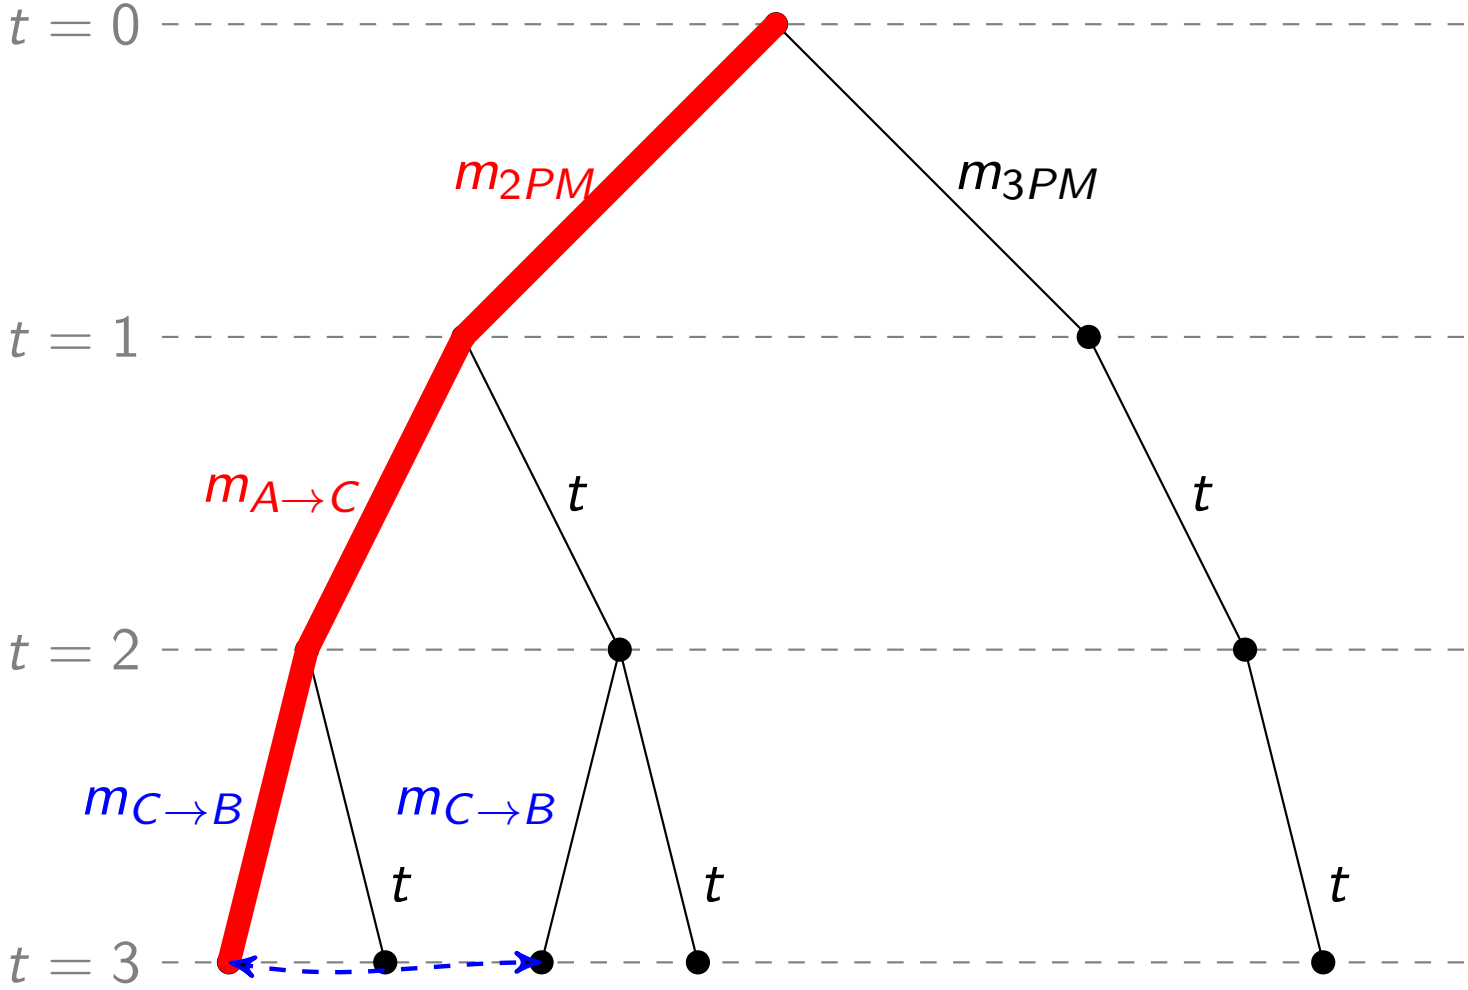

Bob's uncertainty + 'Protocol information':

$$H, 3 \models \neg K_B K_A K_B P_{2PM}$$

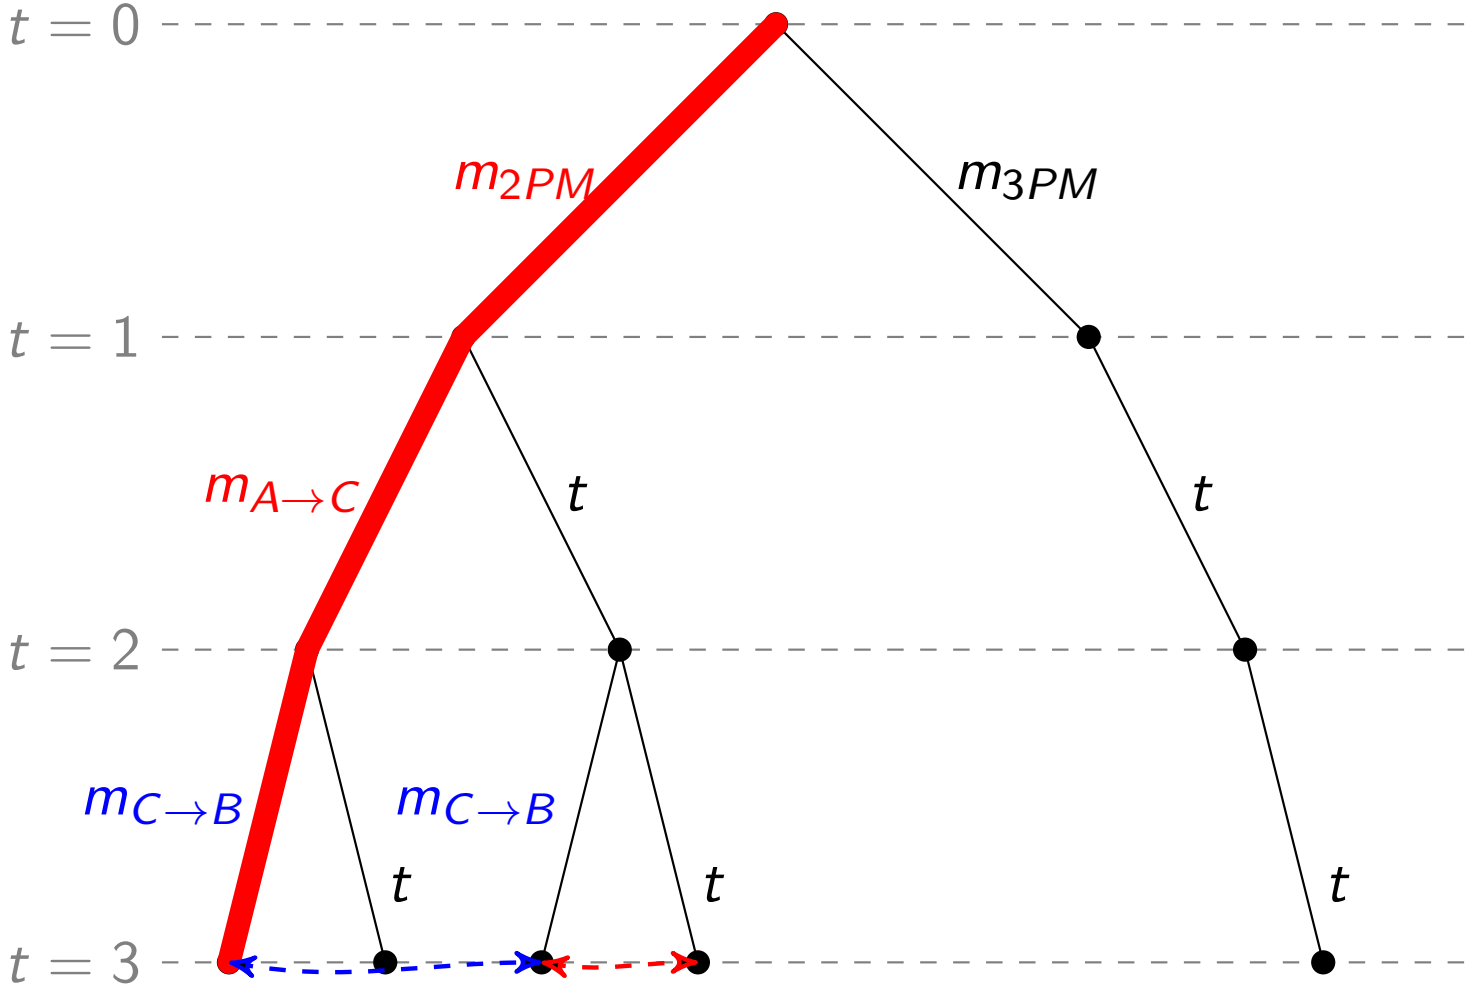

Bob's uncertainty + 'Protocol information':

$$H, 3 \models \neg K_B K_A K_B P_{2PM}$$

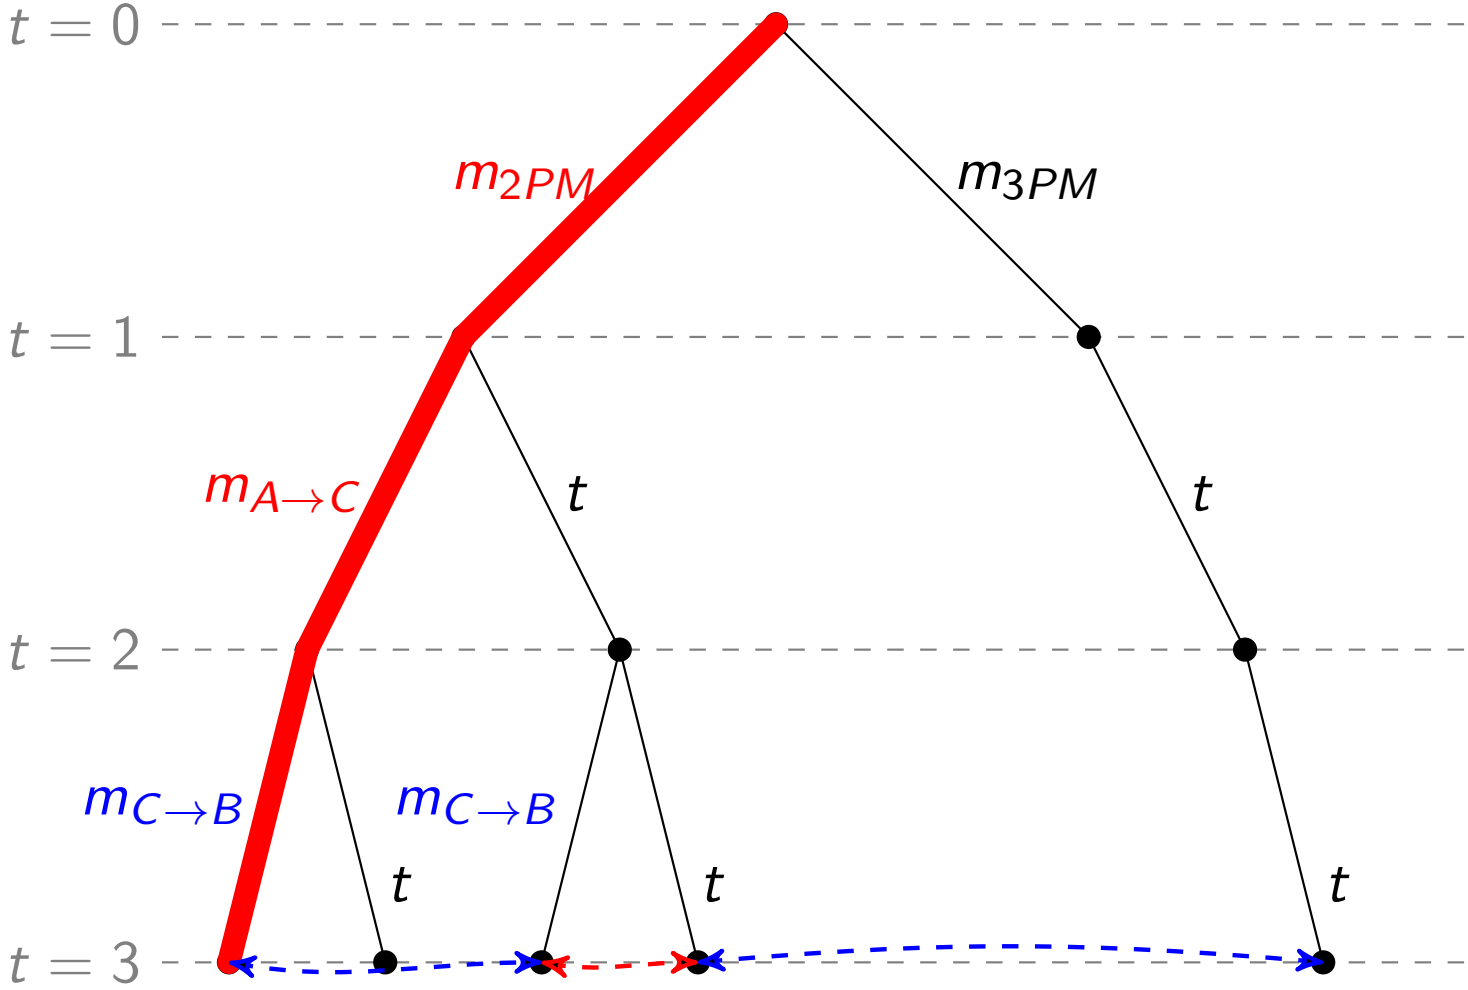

Bob's uncertainty + 'Protocol information':

$$H, 3 \models \neg K_B K_A K_B P_{2PM}$$

Parameters of the Logical Framework

Parameters of the Logical Framework

1. Expressivity of the formal language. Does the language include a common knowledge operator? A future operator? Both?

Parameters of the Logical Framework

1. Expressivity of the formal language. Does the language include a common knowledge operator? A future operator? Both?

2. Structural conditions on the underlying event structure. Do we restrict to protocol frames (finitely branching trees)? Finitely branching forests? Or, arbitrary ETL frames?

Parameters of the Logical Framework

1. Expressivity of the formal language. Does the language include a common knowledge operator? A future operator? Both?

2. Structural conditions on the underlying event structure. Do we restrict to protocol frames (finitely branching trees)? Finitely branching forests? Or, arbitrary ETL frames?

3. Conditions on the reasoning abilities of the agents. Do the agents satisfy perfect recall? No miracles? Do they agents' know what time it is?

Agent Oriented Properties:

- **No Miracles**: For all finite histories $H, H' \in \mathcal{H}$ and events $e \in \Sigma$ such that $He \in \mathcal{H}$ and $H'e \in \mathcal{H}$, if $H \sim_i H'$ then $He \sim_i H'e$.
- **Perfect Recall**: For all finite histories $H, H' \in \mathcal{H}$ and events $e \in \Sigma$ such that $He \in \mathcal{H}$ and $H'e \in \mathcal{H}$, if $He \sim_i H'e$ then $H \sim_i H'$.
- **Synchronous**: For all finite histories $H, H' \in \mathcal{H}$, if $H \sim_i H'$ then $\mathrm{len}(H) = \mathrm{len}(H')$.

Perfect Recall

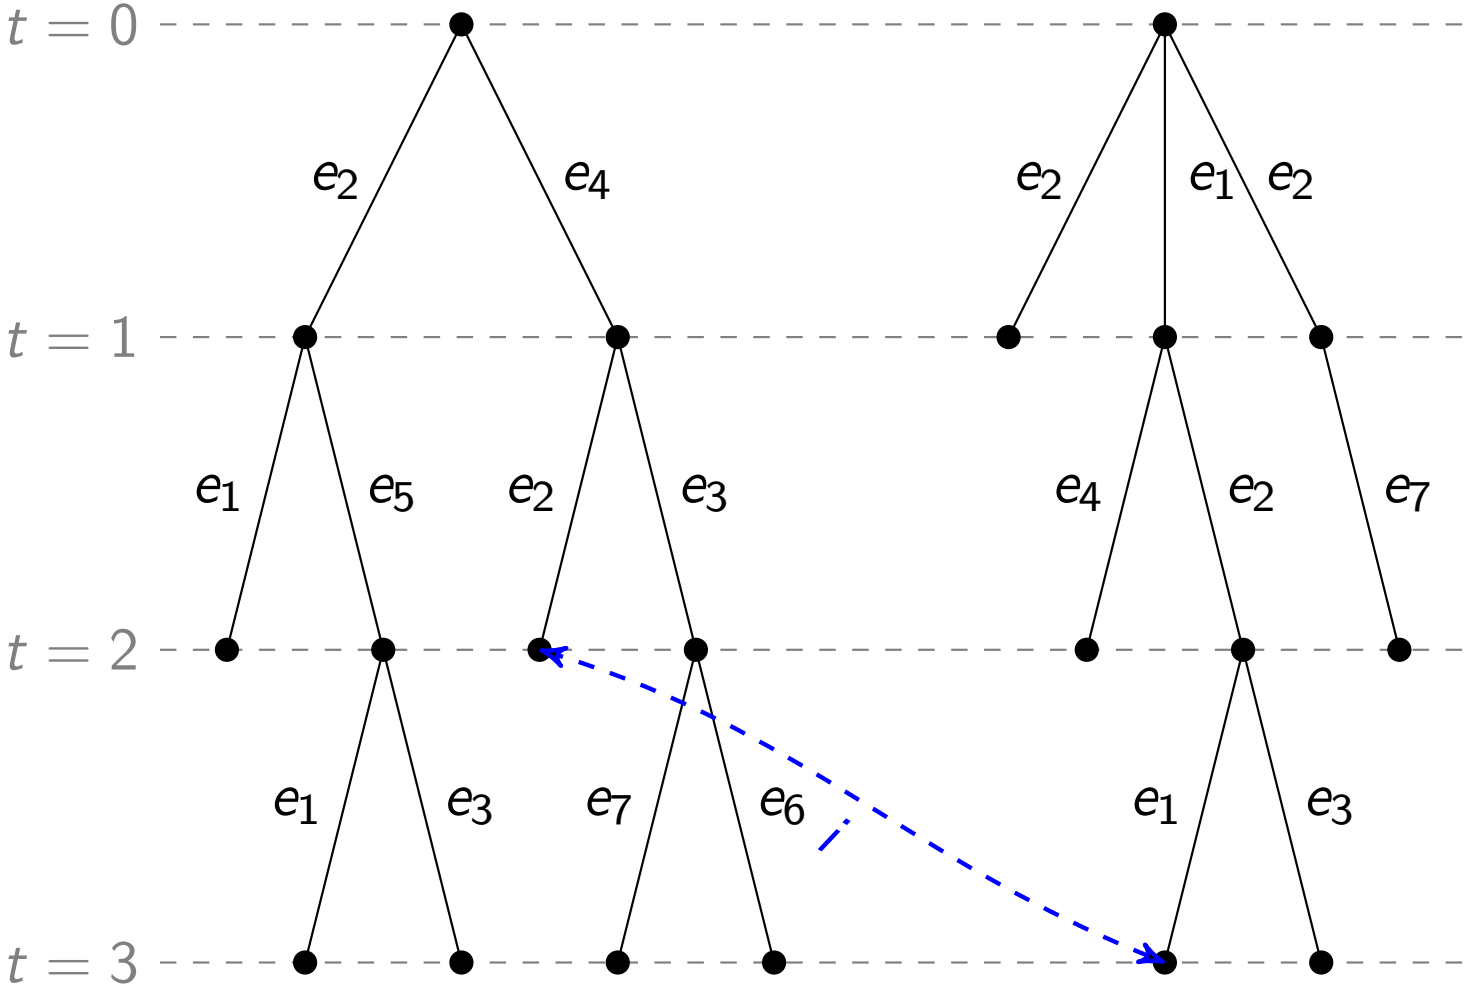

Perfect Recall

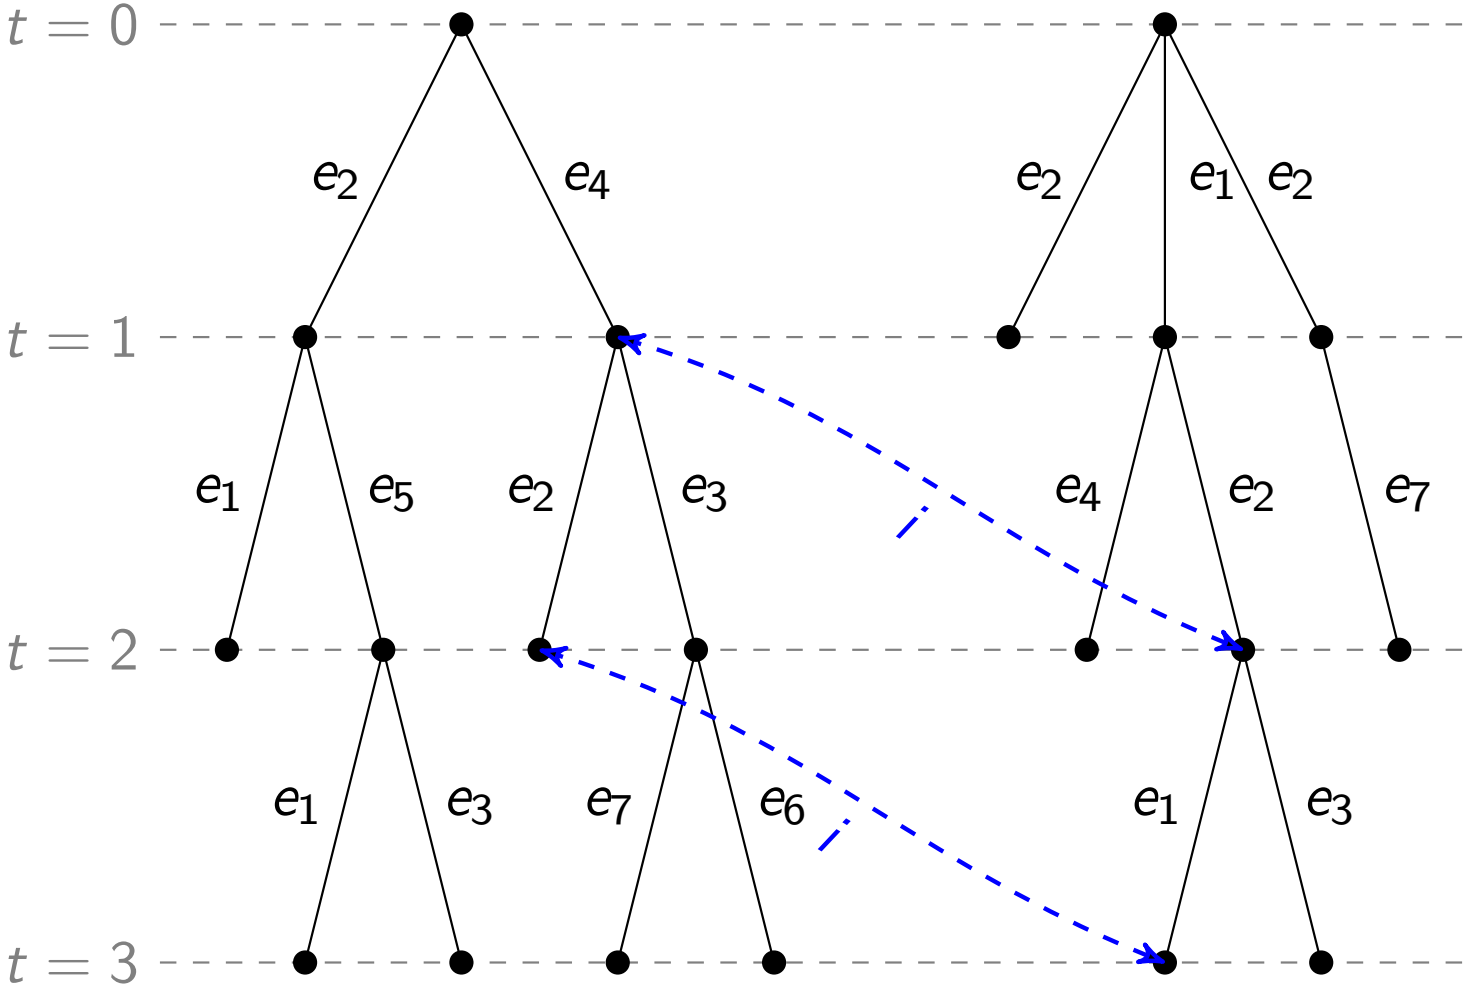

Perfect Recall

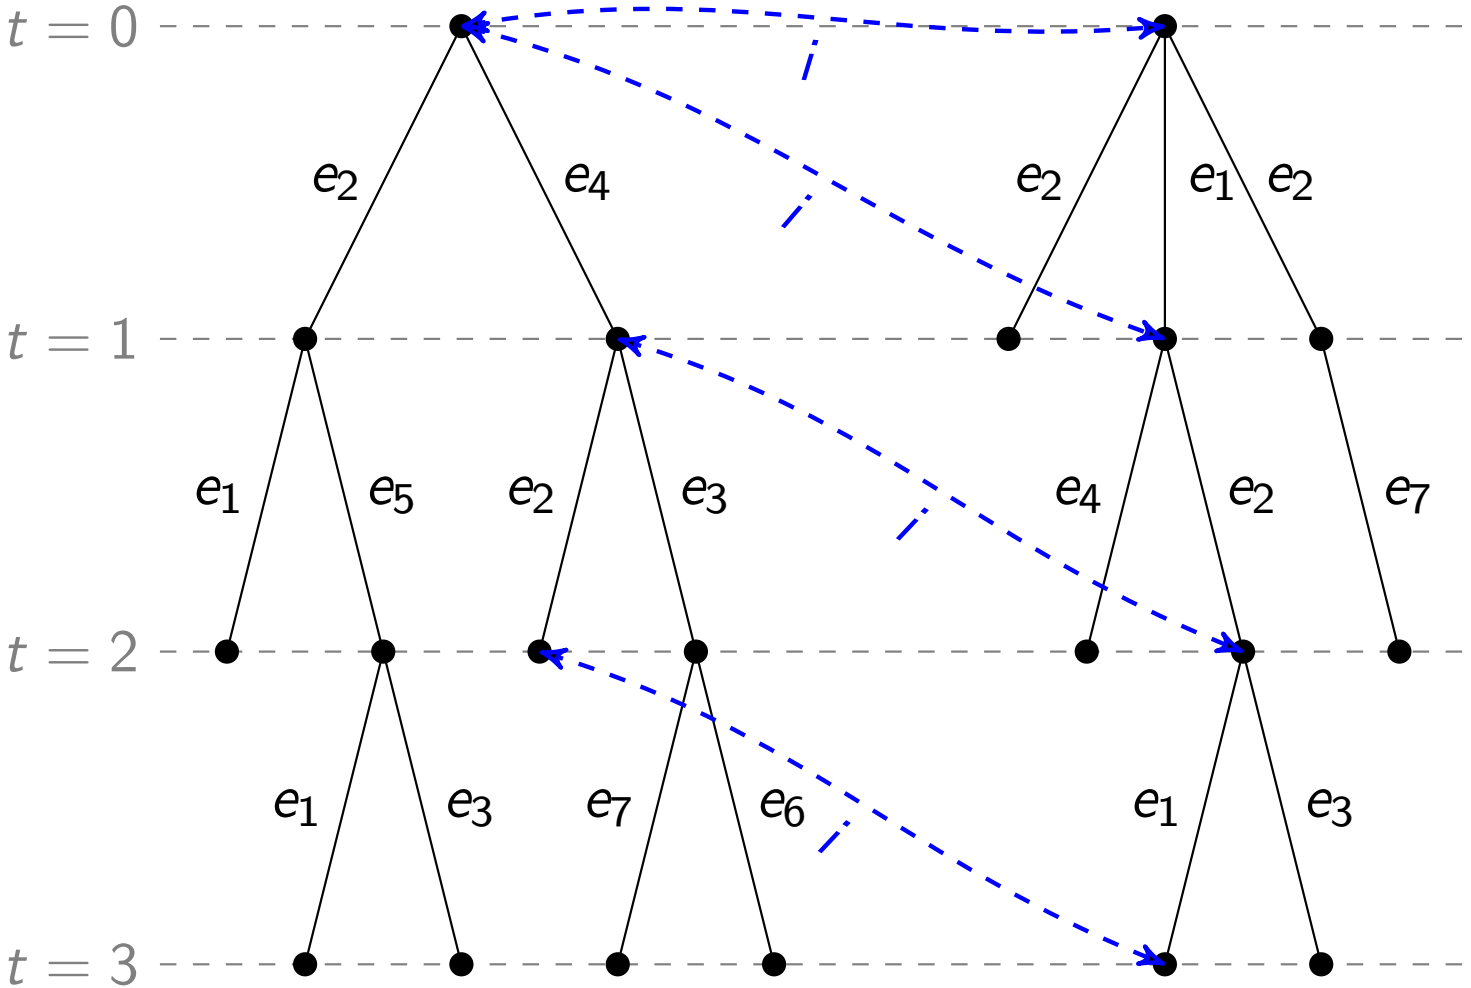

No Miracles

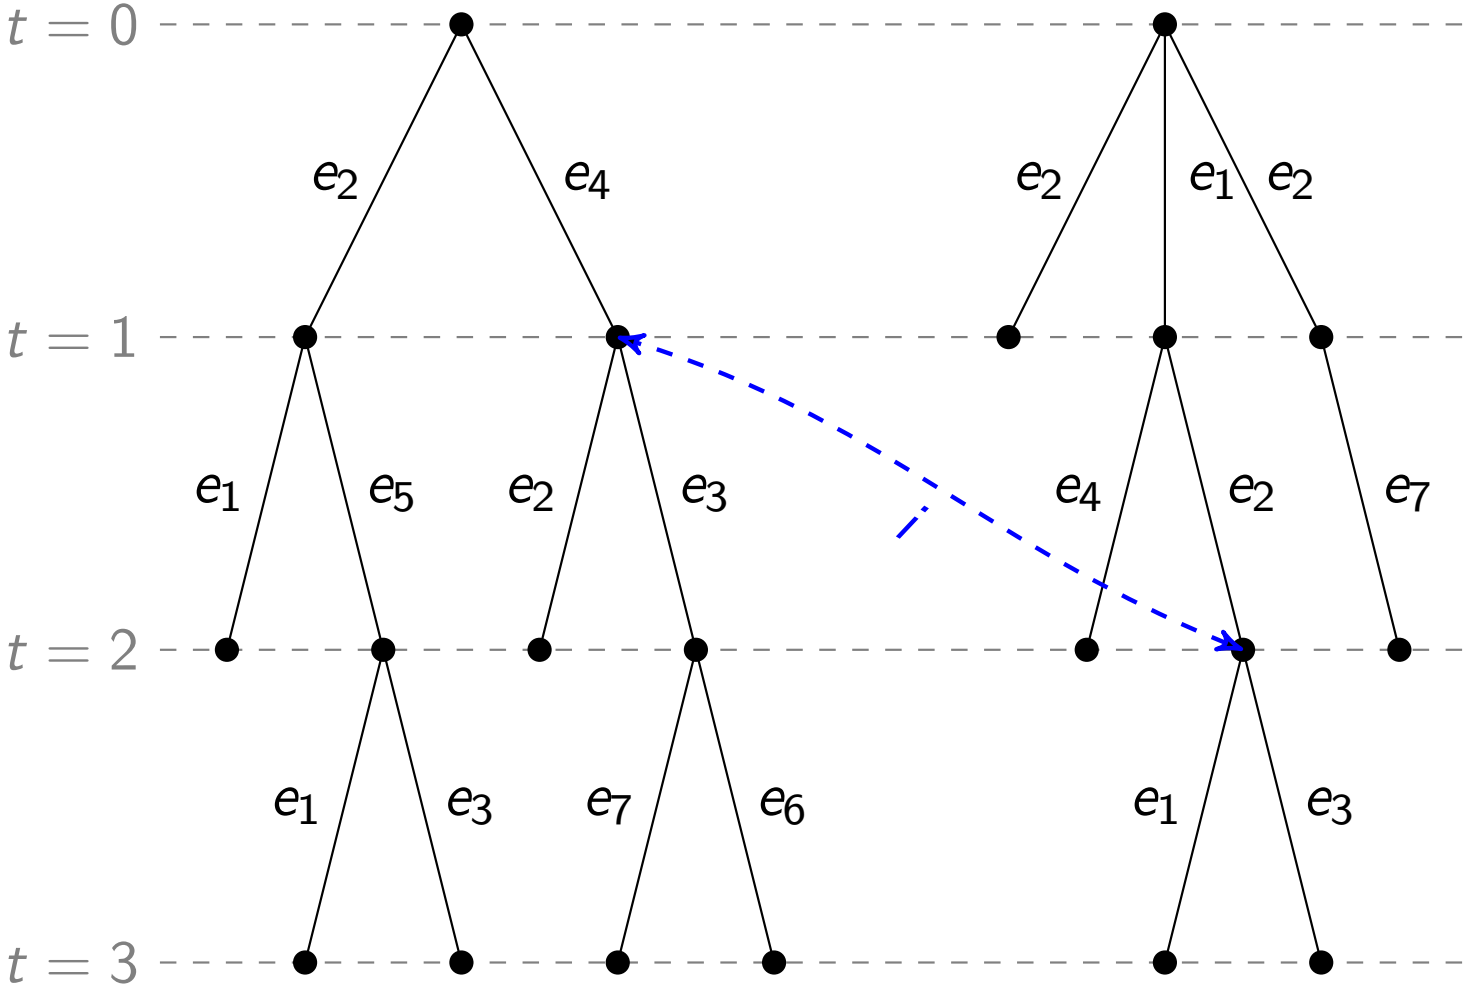

No Miracles

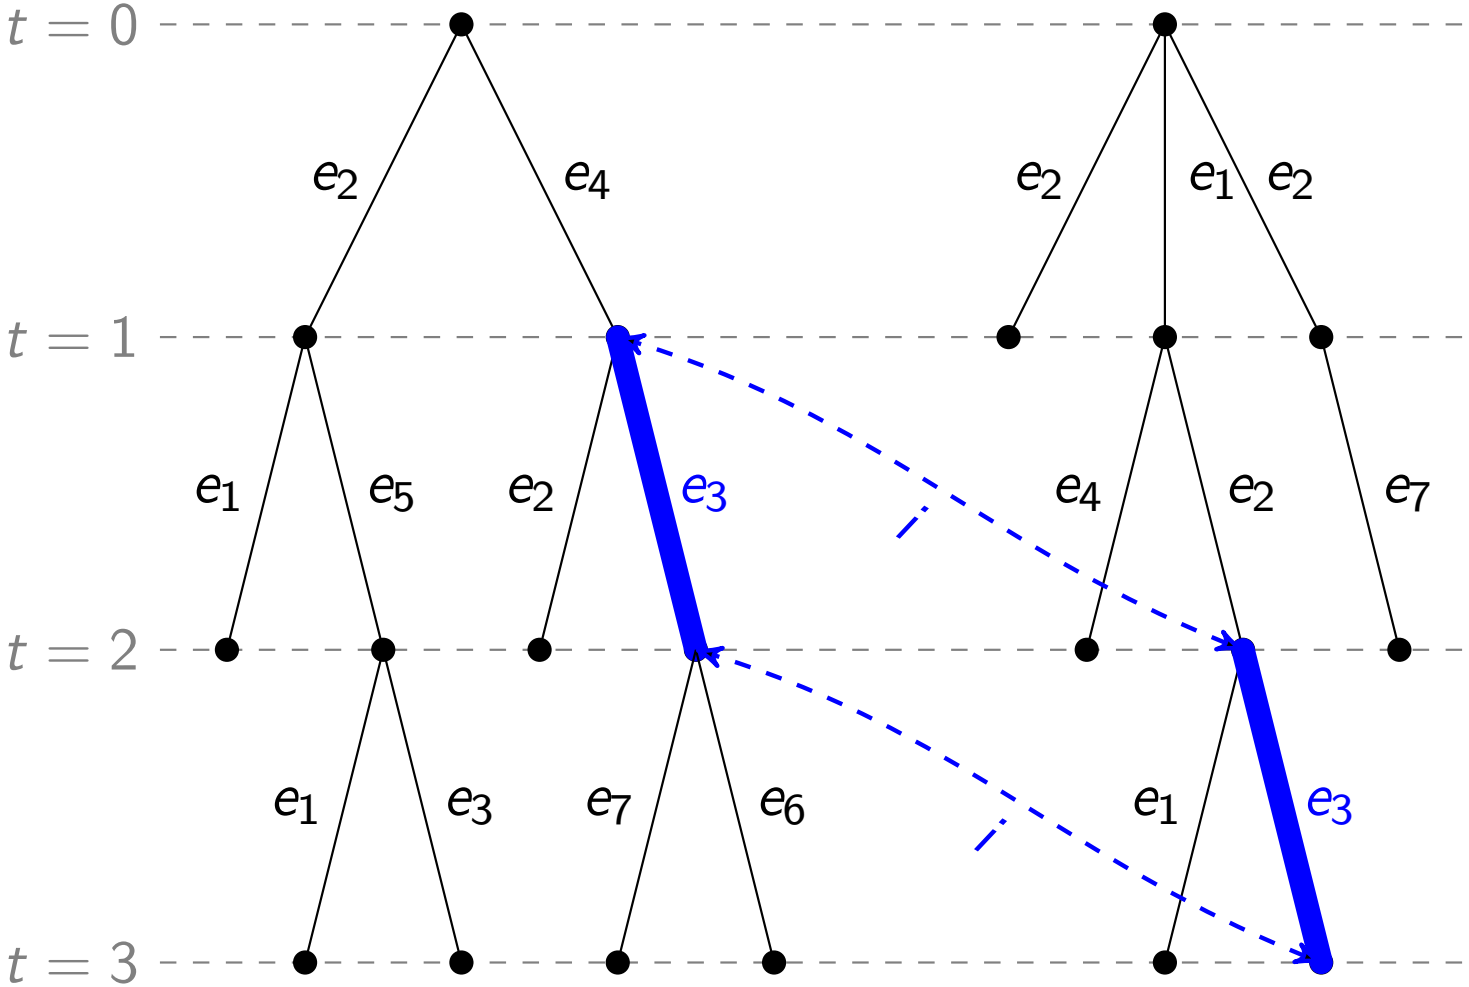

No Miracles

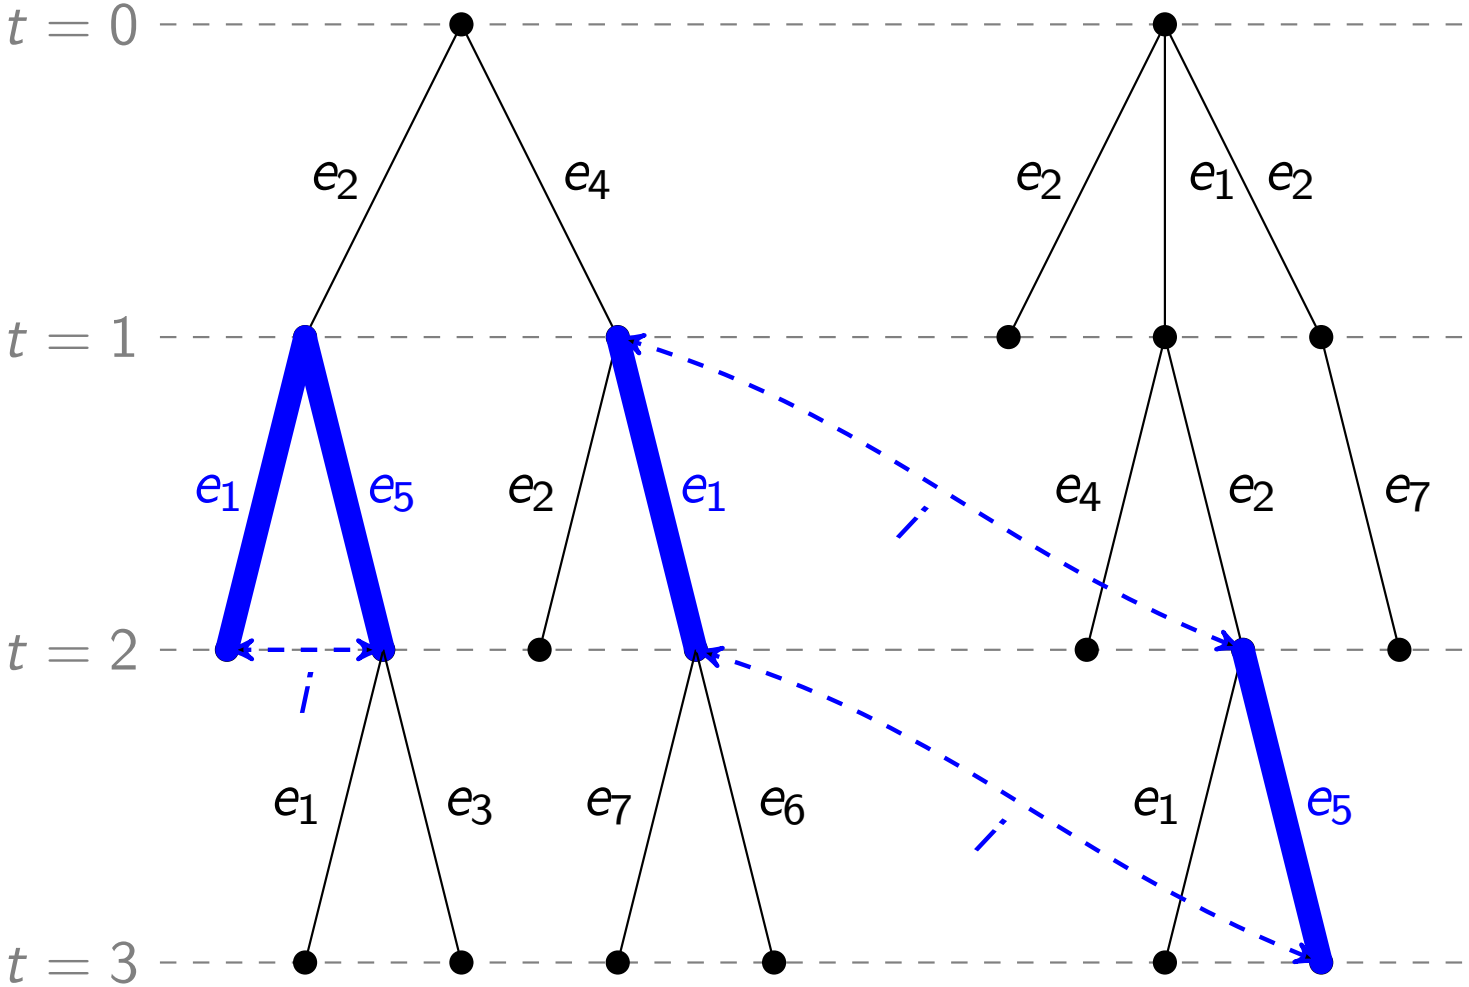

Ideal Agents

Assume there are two agents

Theorem

The logic of ideal *agents with respect to a language with common knowledge and future is highly undecidable (for example, by assuming perfect recall).*

J. Halpern and M. Vardi.. *The Complexity of Reasoning abut Knowledge and Time. J. Computer and Systems Sciences*, 38, 1989.

J. van Benthem and EP. *The Tree of Knowledge in Action*. Proceedings of AiML, 2006.

Constrained Public Announcement

1. $A \to \langle A \rangle \top$ vs. $\langle A \rangle \top \to A$

Constrained Public Announcement

1. $A \rightarrow \langle A \rangle \top$ vs. $\langle A \rangle \top \rightarrow A$

2. $\langle A \rangle K_i P \leftrightarrow A \wedge K_i \langle A \rangle P$

Constrained Public Announcement

1. $A \to \langle A \rangle \top$ vs. $\langle A \rangle \top \to A$
2. $\langle A \rangle K_i P \leftrightarrow A \wedge K_i \langle A \rangle P$
3. $\langle A \rangle K_i P \leftrightarrow \langle A \rangle \top \wedge K_i (A \to \langle A \rangle P)$

Constrained Public Announcement

1. $A \rightarrow \langle A \rangle \top$ vs. $\langle A \rangle \top \rightarrow A$
2. $\langle A \rangle K_i P \leftrightarrow A \wedge K_i \langle A \rangle P$
3. $\langle A \rangle K_i P \leftrightarrow \langle A \rangle \top \wedge K_i (A \rightarrow \langle A \rangle P)$
4. $\langle A \rangle K_i P \leftrightarrow \langle A \rangle \top \wedge K_i (\langle A \rangle \top \rightarrow \langle A \rangle P)$

1. The agents' *observational* powers.

2. The *type* of change triggered by the event.

3. The underlying *protocol* specifying which events (observations, messages, actions) are available (or permitted) at any given moment.

1. The agents' *observational* powers.

2. The *type* of change triggered by the event.

3. The underlying *protocol* specifying which events (observations, messages, actions) are available (or permitted) at any given moment.

Agents may differ in precisely how they incorporate new information into their epistemic states. These differences are based, in part, on the agents' perception of the *source* of the information. For example, an agent may consider a particular source of information *infallible* (not allowing for the possibility that the source is mistaken) or merely *trustworthy* (accepting the information as reliable, though allowing for the possibility of a mistake).

Informative Actions

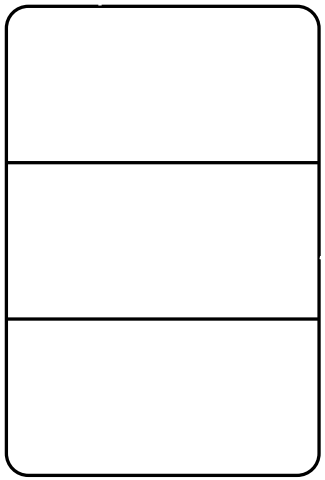

Informative Actions

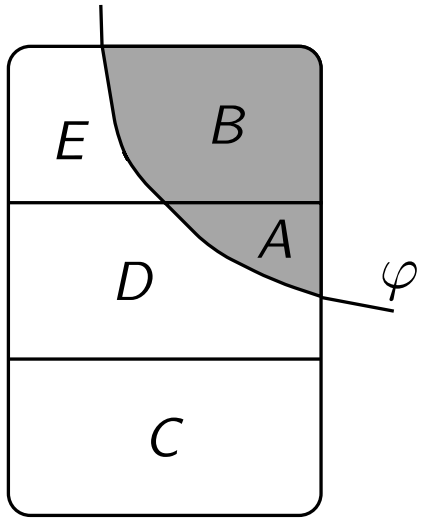

Incorporate the new information φ

Informative Actions

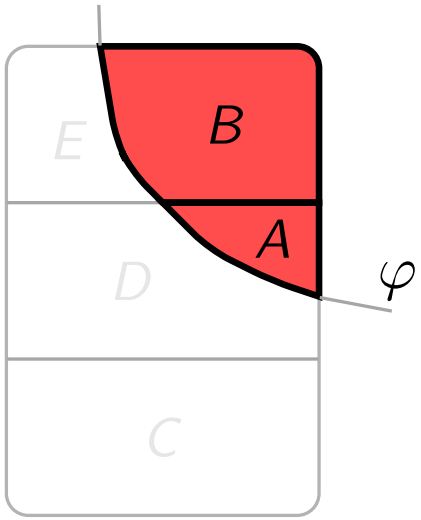

Public Announcement: Information from an infallible source
$(!\varphi)$: $A \prec_i B$

Informative Actions

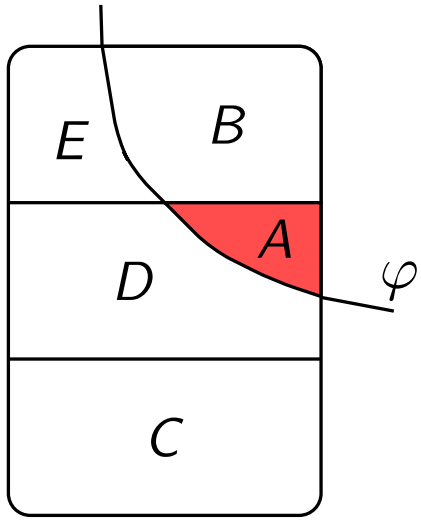

Public Announcement: Information from an infallible source
$(!\varphi)$: $A \prec_i B$

Conservative Upgrade: Information from a trusted source
$(\Uparrow\varphi)$: $A \prec_i C \prec_i D \prec_i B \cup E$

Informative Actions

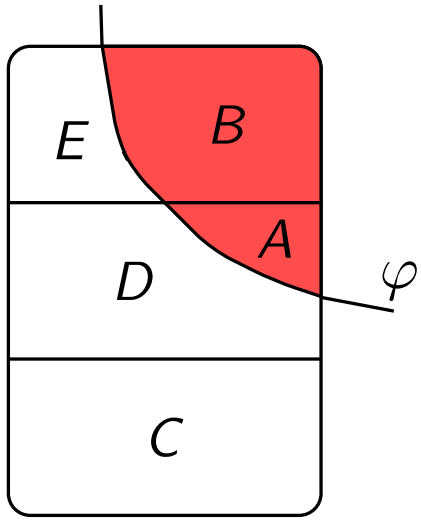

Public Announcement: Information from an infallible source $(!\varphi)$: $A \prec_i B$

Conservative Upgrade: Information from a trusted source $(\uparrow\varphi)$: $A \prec_i C \prec_i D \prec_i B \cup E$

Radical Upgrade: Information from a strongly trusted source $(\Uparrow\varphi)$: $A \prec_i B \prec_i C \prec_i D \prec_i E$