Results in Modal and Dynamic Epistemic Logic

A Formalization in Lean

https://www.github.com/paulaneeley/modal

Paula Neeley
Supervised by Jeremy Avigad

Carnegie Mellon University

Jan 6th, 2021

What is Dynamic Epistemic Logic?

What is Dynamic Epistemic Logic?

Dynamic Epistemic Logic is a logical framework for reasoning about how knowledge changes for systems of agents as communication events occur.

What is Dynamic Epistemic Logic?

Dynamic Epistemic Logic is a logical framework for reasoning about how knowledge changes for systems of agents as communication events occur.

It is useful within a wide variety of disciplines, such as:

- economics
- game theory
- artificial intelligence
- cryptography

What is Dynamic Epistemic Logic?

Dynamic Epistemic Logic is a logical framework for reasoning about how knowledge changes for systems of agents as communication events occur.

It is useful within a wide variety of disciplines, such as:

- economics
- game theory
- artificial intelligence
- cryptography

We will focus on a fragment of Dynamic Epistemic Logic called Public Announcement Logic without common knowledge (PAL).

What is Dynamic Epistemic Logic?

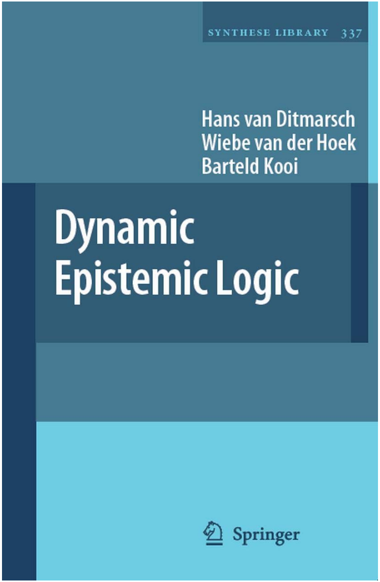

Roadmap

1. Theory of Public Announcement Logic
   - Motivating Example: The Muddy Children Puzzle
   - Dynamic Operators and Frame Restrictions
   - Soundness and Completeness Results

2. Conclusions and Future Work
   - Frame Definability and Undefinability
   - Topological Semantics

The Muddy Children Puzzle

The Muddy Children Puzzle

Three children go to the park to play. When their father comes to find them, he sees that two of them have mud on their foreheads.

The Muddy Children Puzzle

Three children go to the park to play. When their father comes to find them, he sees that two of them have mud on their foreheads.

The father announces, "At least one of you has mud on your forehead", and then asks them, "Do you know if you have mud on your forehead?" The children simultaneously respond, "I don't know."

The Muddy Children Puzzle

Three children go to the park to play. When their father comes to find them, he sees that two of them have mud on their foreheads.

The father announces, "At least one of you has mud on your forehead", and then asks them, "Do you know if you have mud on your forehead?" The children simultaneously respond, "I don't know."

The father then repeats his question, "Do you know if you have mud on your forehead?" This time the two children with muddy foreheads simultaneously answer, "Yes, I do!" while the remaining child again answers, "I don't know."

The Muddy Children Puzzle

Theorem: If there are n children, $k \leq n$ of whom are muddy, then the k children can know that they are muddy after the father repeats his question k times.

Proof: By induction on k.

The Muddy Children Puzzle

Initial states of uncertainty:

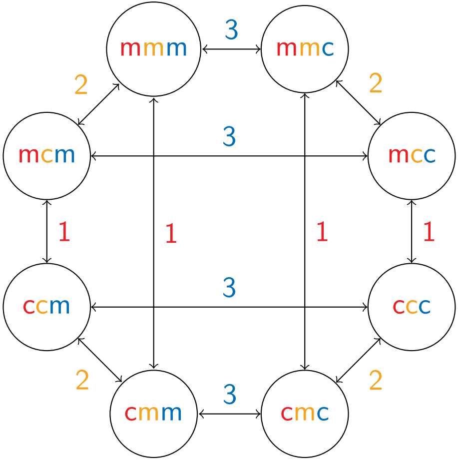

The Muddy Children Puzzle

Initial states of uncertainty:

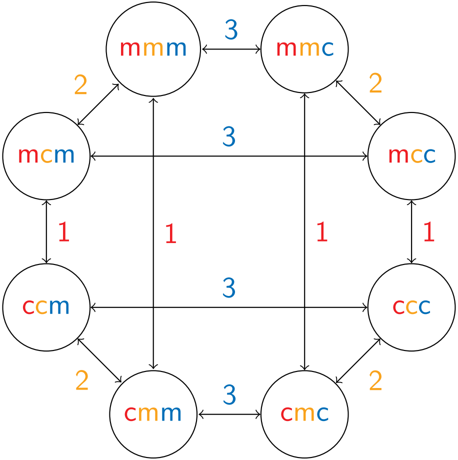

- Nodes represent 2^3 possible worlds (3 children, each of whom are either muddy or not)

The Muddy Children Puzzle

Initial states of uncertainty:

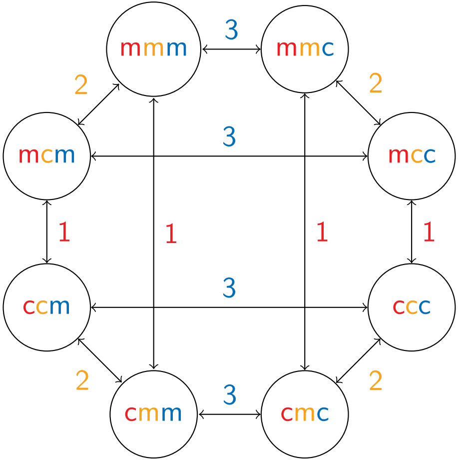

- Nodes represent 2^3 possible worlds (3 children, each of whom are either muddy or not)
- Numbered edges represent each child's uncertainty between states

The Muddy Children Puzzle

Initial states of uncertainty:

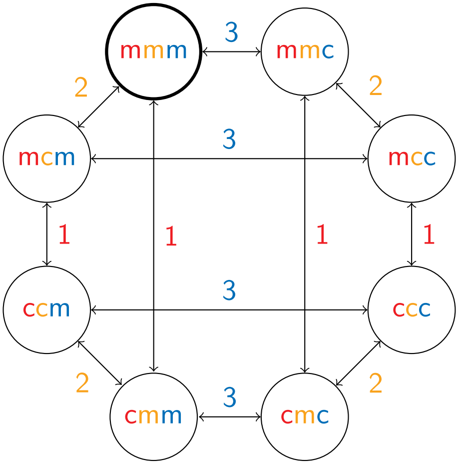

- Nodes represent 2^3 possible worlds (3 children, each of whom are either muddy or not)
- Numbered edges represent each child's uncertainty between states

The Muddy Children Puzzle

Initial states of uncertainty:

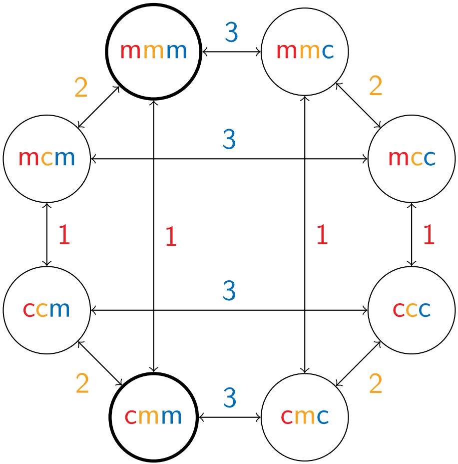

- Nodes represent 2^3 possible worlds (3 children, each of whom are either muddy or not)
- Numbered edges represent each child's uncertainty between states

The Muddy Children Puzzle

States of uncertainty after the father announces, “At least one of you has mud on your forehead”:

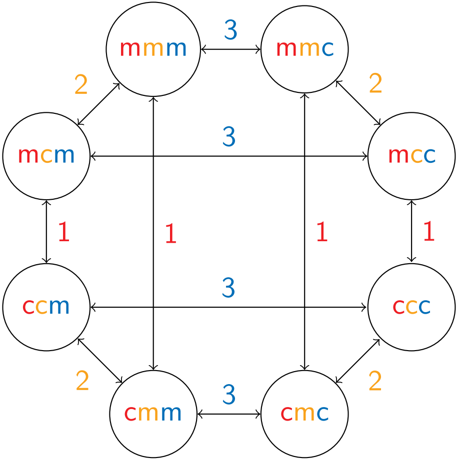

The Muddy Children Puzzle

States of uncertainty after the father announces, "At least one of you has mud on your forehead":

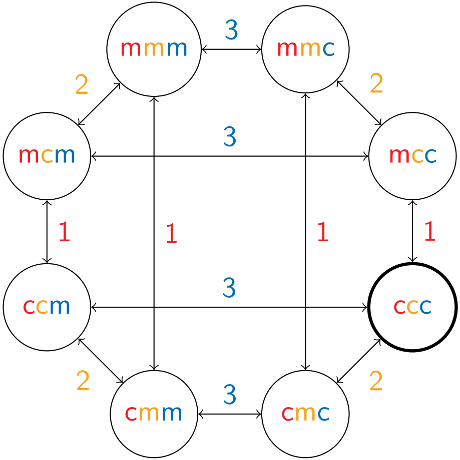

The Muddy Children Puzzle

States of uncertainty after the father announces, "At least one of you has mud on your forehead":

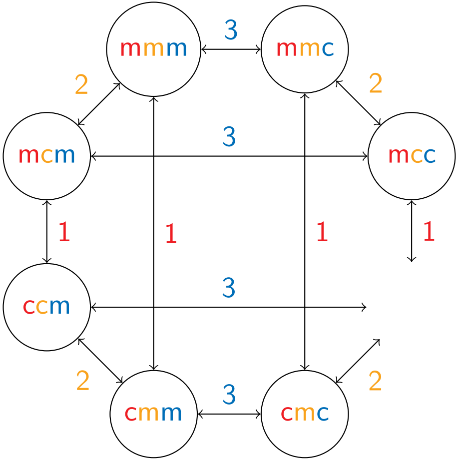

The Muddy Children Puzzle

States of uncertainty after the father announces, "At least one of you has mud on your forehead":

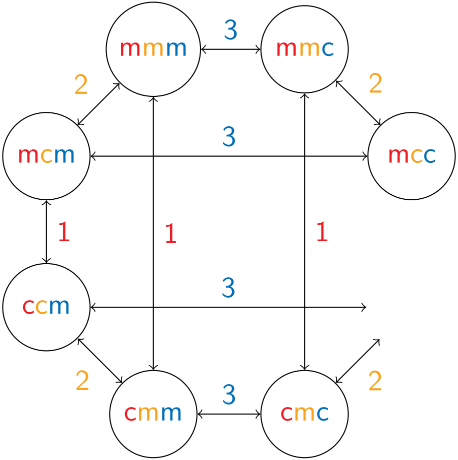

- At mcc, child 1 is certain mcc is the correct state,

The Muddy Children Puzzle

States of uncertainty after the father announces, "At least one of you has mud on your forehead":

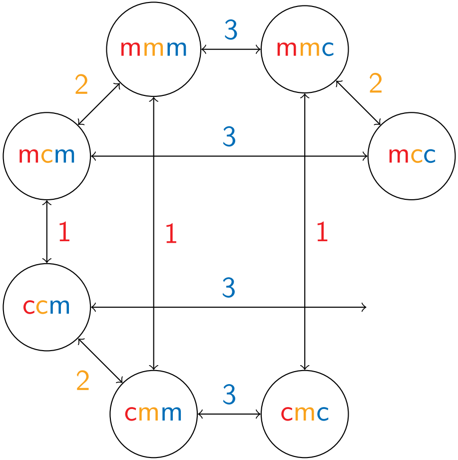

- At mcc, child 1 is certain mcc is the correct state,
- At cmc, child 2 is certain cmc is the correct state,

The Muddy Children Puzzle

States of uncertainty after the father announces, "At least one of you has mud on your forehead":

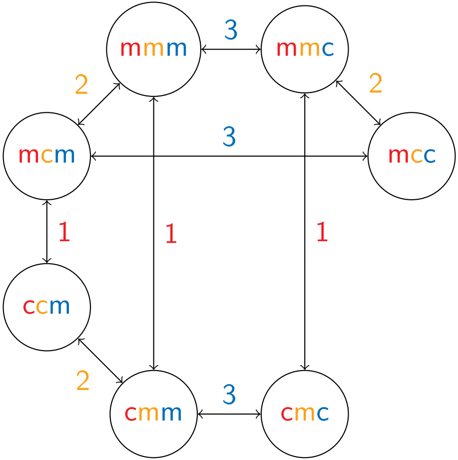

- At mcc, child 1 is certain mcc is the correct state,
- At cmc, child 2 is certain cmc is the correct state,
- At ccm, child 3 is certain ccm is the correct state.

The Muddy Children Puzzle

States of uncertainty after the children simultaneously announce, "I don't know":

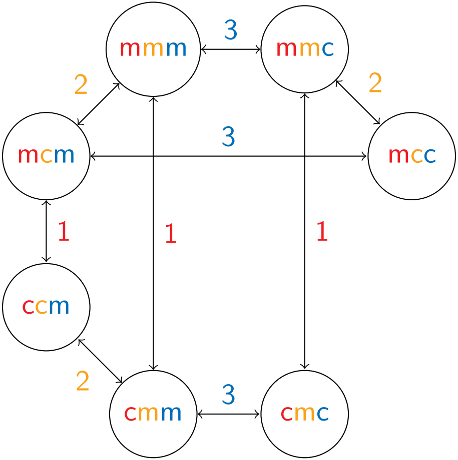

The Muddy Children Puzzle

States of uncertainty after the children simultaneously announce, "I don't know":

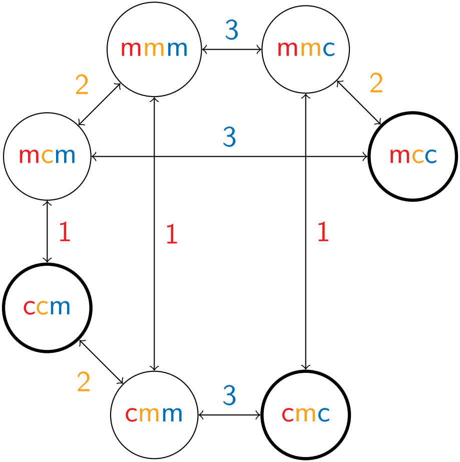

The Muddy Children Puzzle

States of uncertainty after the children simultaneously announce, “I don’t know”:

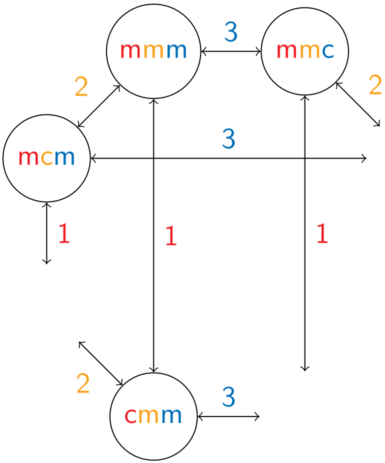

The Muddy Children Puzzle

States of uncertainty after the children simultaneously announce, "I don't know":

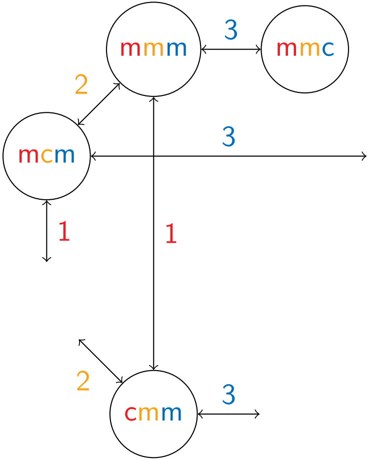

- At mmc, children 1 and 2 are certain mmc is the correct state,

The Muddy Children Puzzle

States of uncertainty after the children simultaneously announce, "I don't know":

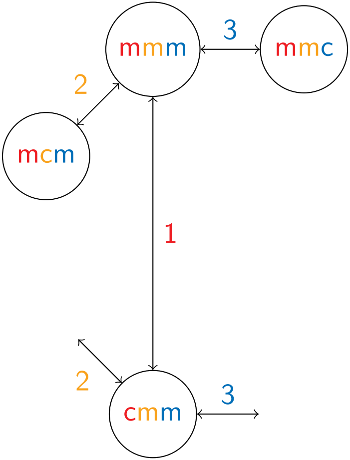

- At mmc, children 1 and 2 are certain mmc is the correct state,
- At mcm, children 1 and 3 are certain mcm is the correct state,

The Muddy Children Puzzle

States of uncertainty after the children simultaneously announce, "I don't know":

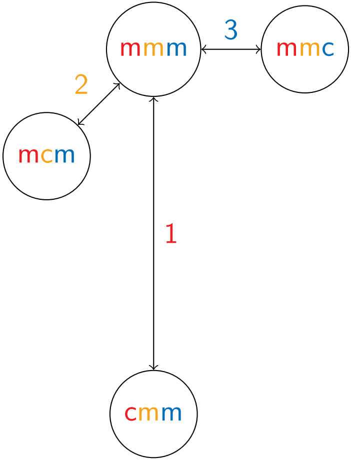

- At mmc, children 1 and 2 are certain mmc is the correct state,
- At mcm, children 1 and 3 are certain mcm is the correct state,
- At cmm, children 2 and 3 are certain cmm is the correct state.

The Muddy Children Puzzle

States of uncertainty after the father repeats his question, "Do you know if you have mud on your forehead?"

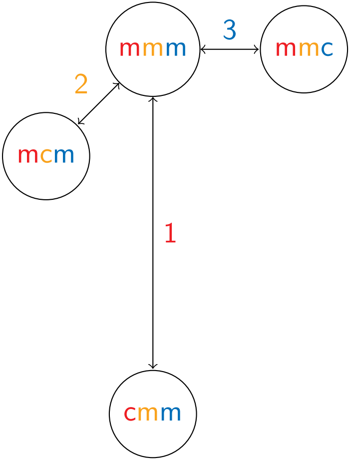

The Muddy Children Puzzle

States of uncertainty after the father repeats his question, “Do you know if you have mud on your forehead?”

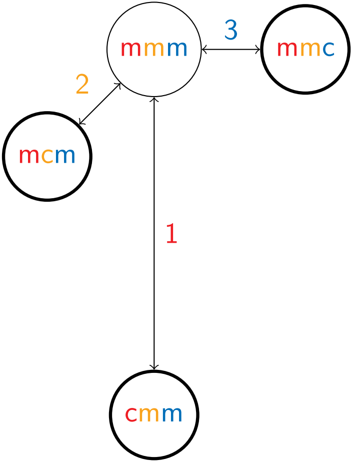

This time the two children with muddy foreheads simultaneously answer, “Yes, I do!” while the remaining child answers, “I don’t know.”

Theory of Public Announcement Logic

Dynamic Operators and Frame Restrictions
Soundness and Completeness Results

Public Announcement Logic

Definition (The Language of Epistemic Logic)

Given a finite set of agents A and a countable set of primitive propositions $\textsc{Prop}$, the language $\mathcal{L}_K$ is defined inductively as follows:

$$\phi := \bot \mid p_n \mid \phi \to \psi \mid K_a \phi$$

where $a \in A$ and $p_n \in \textsc{Prop}$.

The necessity operator $K_a\phi$ is read as "agent a knows that ϕ."

(This is just the basic modal language indexed over agents.)

Public Announcement Logic

Definition (The Language of Public Announcement Logic)

Given a finite set of agents A and a countable set of primitive propositions $\textsc{Prop}$, the language $\mathcal{L}_{K[]}$ is defined inductively as follows:

$$\phi := \bot \mid p_n \mid \phi \to \psi \mid K_a\phi \mid [\phi]\psi$$

where $a \in A$ and $p_n \in \textsc{Prop}$.

The update operator $[\phi]\psi$ is read as, "after every truthful announcement of ϕ, ψ holds."

Public Announcement Logic

Definition (Kripke Frame)

A Kripke frame is a tuple $F = (W, R^A)$ where W denotes a non-empty set of possible worlds and R^A is a function, yielding for each $a \in A$ a binary relation $R_a \subseteq W \times W$ (called the *accessibility relation*) between worlds.

wR_ay means that world y is accessible from world w for agent a.

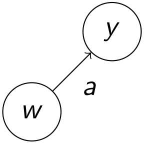

Public Announcement Logic

Definition (Kripke Model)

A Kripke model $M = ((W, R^A), v)$ is a tuple where (W, R^A) denotes a Kripke frame and $v : \text{PROP} \to 2^W$ is a valuation function over the primitive propositions.

$w \in v(p_n)$ means that "p_n is true at world w,"

$[\![p_n]\!]_M$ denotes the set of all worlds in M where p_n is true.

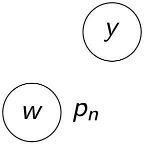

Public Announcement Logic

Definition (Validity for Epistemic Logic)

Let w be a world in a model $M = ((W, R^A), v)$. We say that ϕ is valid for agent a at world w in M whenever

$$
\begin{array}{lll}
(M, w) \models \bot & \text{iff} & \text{false} \\
(M, w) \models p_n & \text{iff} & w \in v(p_n) \\
(M, w) \models \phi \to \psi & \text{iff} & (M, w) \models \phi \quad \text{implies } (M, w) \models \psi \\
(M, w) \models K_a\phi & \text{iff} & \forall y \in W,\ wR_a y \text{ implies } (M, y) \models \phi
\end{array}
$$

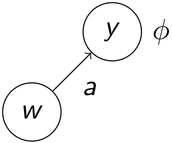

Public Announcement Logic

Definition (Validity for Public Announcement Logic)

Let w be a world in a model $M = ((W, R^A), v)$. We say that ϕ is valid for agent a at world w in M whenever

$$
\begin{array}{lll}
(M, w) \models \bot & \text{iff} & \text{false} \\
(M, w) \models p_n & \text{iff} & w \in v(p_n) \\
(M, w) \models \phi \to \psi & \text{iff} & (M, w) \models \phi \quad \text{implies } (M, w) \models \psi \\
(M, w) \models K_a\phi & \text{iff} & \forall y \in W,\ wR_a y \text{ implies } (M, y) \models \phi \\
(M, w) \models [\phi]\psi & \text{iff} & (M, w) \models \phi \quad \text{implies } (M|\phi, w) \models \psi
\end{array}
$$

Where $M|\phi = ((W', R^{A\prime}), v')$ is such that,

$$
\begin{aligned}
W' &= [\![\phi]\!]_M \\
R^{A\prime} &= R^A \cap ([\![\phi]\!]_M \times [\![\phi]\!]_M) \\
v'(p_n) &= v(p_n) \cap [\![\phi]\!]_M
\end{aligned}
$$

Public Announcement Logic

Definition (Proof System S5 for Epistemic Logic)

Given a set of agents A and primitive propositions PROP, the axiomatic system $S5$ is comprised of:

- all instances of propositional tautologies,
- all instances of the schemes:

  $K_a(\phi \to \psi) \to (K_a\phi \to K_a\psi)$, (DISTRIBUTION OF K_a OVER $\to$)

  $K_a\phi \to \phi$, (TRUTH)

  $K_a\phi \to K_aK_a\phi$, (POSITIVE INTROSPECTION)

  $\neg K_a\phi \to K_a\neg K_a\phi$, (NEGATIVE INTROSPECTION)

- from ϕ and $\phi \to \psi$, we infer ψ, (MODUS PONENS)
- from ϕ, we infer $K_a\phi$. (NECESSITATION)

Public Announcement Logic

Definition (Proof System PA for Public Announcement Logic)

Given a set of agents A and primitive propositions $\textsc{prop}$, the axiomatic system PA is comprised of:

- all instances of propositional tautologies,
- all instances of the schemes:

$K_a(\phi \to \psi) \to (K_a\phi \to K_a\psi)$,	(DISTRIBUTION OF K_a OVER $\to$)
$K_a\phi \to \phi$,	(TRUTH)
$K_a\phi \to K_aK_a\phi$,	(POSITIVE INTROSPECTION)
$\neg K_a\phi \to K_a\neg K_a\phi$,	(NEGATIVE INTROSPECTION)
$[\phi]\bot \leftrightarrow (\phi \to \bot)$,	(ATOMIC FALSEHOOD)
$[\phi]p_n \leftrightarrow (\phi \to p_n)$,	(ATOMIC PERMANENCE)
$[\phi](\psi \to \chi) \leftrightarrow ([\phi]\psi \to [\phi]\chi)$,	(ANNOUNCEMENT AND IMPLICATION)
$[\phi]K_a\psi \leftrightarrow (\phi \to K_a[\phi]\psi)$,	(ANNOUNCEMENT AND KNOWLEDGE)
$[\phi][\psi]\chi \leftrightarrow [\phi \wedge [\phi]\psi]\chi$,	(ANNOUNCEMENT COMPOSITION)

- from ϕ and $\phi \to \psi$, we infer ψ, (MODUS PONENS)
- from ϕ, we infer $K_a\phi$. (NECESSITATION)

Public Announcement Logic

Definition (Provability)

Given a proof system AX ($S5$ or PA), a formula ϕ is provable from AX if either:

- ϕ is an axiom,
- ϕ follows from provable formulas ψ and χ by modus ponens,
- ϕ follows from provable formula ψ by necessitation.

We denote the fact that ϕ is provable by $\vdash_{AX} \phi$.

Public Announcement Logic

Completeness for Epistemic Logic

Theorem (Completeness)

The proof system $S5$ is sound and complete with respect to the class of frames $\mathcal{F}_{eq}$ whose relation is an equivalence relation (reflexive, symmetric, transitive). That is,

$$\vdash_{S5} \phi \iff \mathcal{F}_{eq} \models \phi.$$

Proof.

($\Rightarrow$) By induction on the provability relation.
($\Leftarrow$) By contraposition using a canonical model construction. $\square$

Theorem (Completeness)

The proof system *PA* is sound and complete with respect to the class of frames $\mathcal{F}_{eq}$ whose relation is an equivalence relation (reflexive, symmetric, transitive). That is,

$$\vdash_{PA} \phi \iff \mathcal{F}_{eq} \models \phi.$$

Proof.

($\Rightarrow$) By induction on the provability relation. □

The language $\mathcal{L}_{K[]}$ has the same expressive power as the language $\mathcal{L}_K$.

Proof. (Cont.)

($\Leftarrow$) To prove completeness, it suffices to define a translation $t : \mathcal{L}_{K[]} \to \mathcal{L}_K$ and show that every formula is provably equivalent to its translation. That is,

$$\vdash \phi \leftrightarrow t(\phi).$$

$\square$

Definition (Translation Function)

The translation $t : \mathcal{L}_{K[]} \to \mathcal{L}_K$ is defined recursively as follows:

$$\begin{aligned}
t(\bot) &= \bot \\
t(p_n) &= p_n \\
t(\phi \to \psi) &= t(\phi) \to t(\psi) \\
t(K_a\phi) &= K_a t(\phi) \\
t([\phi]\bot) &= t(\phi \to \bot) \\
t([\phi]p_n) &= t(\phi \to p_n) \\
t([\phi](\psi \to \chi)) &= t([\phi]\psi \to [\phi]\chi) \\
t([\phi]K_a\psi) &= t(\phi \to K_a[\phi]\psi) \\
t([\phi][\psi]\chi) &= t([\phi \wedge [\phi]\psi]\chi)
\end{aligned}$$

Public Announcement Logic

Completeness for Public Announcement Logic

Definition (Translation Function)

The translation $t : \mathcal{L}_{K[]} \to \mathcal{L}_K$ is defined recursively as follows:

$$
\begin{aligned}
t(\bot) &= \bot \\
t(p_n) &= p_n \\
t(\phi \to \psi) &= t(\phi) \to t(\psi) \\
t(K_a\phi) &= K_a t(\phi) \\
t([\phi]\bot) &= t(\phi \to \bot) \\
t([\phi]p_n) &= t(\phi \to p_n) \\
t([\phi](\psi \to \chi)) &= t([\phi]\psi \to [\phi]\chi) \\
t([\phi]K_a\psi) &= t(\phi \to K_a[\phi]\psi) \\
t([\phi][\psi]\chi) &= t([\phi \wedge [\phi]\psi]\chi)
\end{aligned}
$$

Definition (Translation Function)

The translation $t : \mathcal{L}_{K[]} \to \mathcal{L}_K$ is defined recursively as follows:

$$
\begin{aligned}
t(\bot) &= \bot \\
t(p_n) &= p_n \\
t(\phi \to \psi) &= t(\phi) \to t(\psi) \\
t(K_a \phi) &= K_a t(\phi) \\
t([\phi]\bot) &= t(\phi \to \bot) \\
t([\phi]p_n) &= t(\phi \to p_n) \\
t([\phi](\psi \to \chi)) &= t([\phi]\psi \to [\phi]\chi) \\
t([\phi]K_a\psi) &= t(\phi \to K_a[\phi]\psi) \\
t([\phi][\psi]\chi) &= t([\phi \land [\phi]\psi]\chi)
\end{aligned}
$$

Public Announcement Logic

Completeness for Public Announcement Logic

Lemma (Equivalence Under Translation)

For all formulas $\phi \in \mathcal{L}_{K[]}$ it is the case that

$$\vdash \phi \leftrightarrow t(\phi).$$

Since, for example, $[\phi \wedge [\phi]\psi]\chi$ is not a subformula of $[\phi][\psi]\chi$, we must define a complexity measure to prove equivalence of formulas.

Proof.

By induction on the complexity of ϕ. □

Corollary

PA completeness follows from *S*5 completeness.

Conclusions and Future Work

Conclusions

Frame Definability and Undefinability

Definition (Definability)

We say ϕ defines a class of frames $\mathcal{F}$ if, for all frames F,

$$\mathsf{F} \in \mathcal{F} \iff F \models \phi.$$

Conclusions

Frame Definability and Undefinability

Definition (Definability)

We say ϕ defines a class of frames $\mathcal{F}$ if, for all frames F,

$$\mathsf{F} \in \mathcal{F} \iff F \models \phi.$$

Examples:

- $\Box p \to p$ defines the class of reflexive frames,
- $\Box(\Box p \to p) \to \Box p$ (aka Löb's formula) defines the class of frames that are transitive and admit no infinite R-paths.

Conclusions

Frame Definability and Undefinability

Definition (Undefinability)

We say a class of frames $\mathcal{F}$ is undefinable if there is no formula ϕ that defines $\mathcal{F}$.

Definition (Undefinability)

We say a class of frames $\mathcal{F}$ is undefinable if there is no formula ϕ that defines $\mathcal{F}$.

To prove undefinability, we rely on invariance under:

- Disjoint unions
- Generated subframes
- Bisimulations
- Surjective bounded morphisms

Conclusions

Frame Definability and Undefinability

Definition (Undefinability)

We say a class of frames $\mathcal{F}$ is undefinable if there is no formula ϕ that defines $\mathcal{F}$.

To prove undefinability, we rely on invariance under:

- Disjoint unions
- Generated subframes
- Bisimulations
- Surjective bounded morphisms

These results are included in my Lean formalization.

Future Work

Completeness for PAL with Common Knowledge

Definition (The Language of PAL with Common Knowledge)

Given a finite set of agents A and a countable set of primitive propositions $\textsc{prop}$, the language $\mathcal{L}_{KC[]}$ is defined inductively as follows:

$$\psi := \bot \mid p_n \mid \psi \supset \phi \mid K_a\phi \mid C_B\phi \mid [\phi]\psi$$

where $a \in A$, $p_n \in \textsc{prop}$, and $B \subseteq A$.

Since PAL with common knowledge *is* more expressive than the basic modal language, the completeness proof is a bit more involved...

Future Work

Topological Semantics

A topological model is a topological space $(X, \mathcal{T})$ together with a valuation $v : \text{PROP} \to 2^X$.

Definition (Topological Validity)

Truth in a model $M = ((X, \mathcal{T}), v)$ is defined as

$$
\begin{aligned}
[\![p_n]\!]_M &= v(p_n) \\
[\![\neg\phi]\!]_M &= X \setminus [\![\phi]\!]_M \\
[\![\phi \to \psi]\!]_M &= [\![\phi]\!]_M \subseteq [\![\psi]\!]_M \\
[\![\Box\phi]\!]_M &= int([\![\phi]\!]_M)
\end{aligned}
$$

where $int([\![\phi]\!]_M)$ denotes the topological interior of the set $[\![\phi]\!]$.

Thanks for listening!

Thanks for listening!

Questions?