Introduction to Formal Epistemology

Lecture 2a

Eric Pacuit and Rohit Parikh

August 14, 2007

✓ Introduction, Motivation and Basic Epistemic Logic

Lecture 2: Other models of Knowledge, Knowledge in Groups and Group Knowledge

Lecture 3: Reasoning about Knowledge and ....

Lecture 4: Logical Omniscience and Other Problems

Lecture 5: Reasoning about Knowledge in the Context of Social Software

Epistemic Logic

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

Kripke Models: $\mathcal{M} = \langle W, R, V\rangle$ with R an equivalence relation

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 $\mathcal{M}, w \models \psi$
- $\mathcal{M}, w \models K\varphi$ if for each $v \in W$, if wRv, then $\mathcal{M}, v \models \varphi$

Epistemic Logic

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

Kripke Models: $\mathcal{M} = \langle W, R, V\rangle$ with R an equivalence relation

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 $\mathcal{M}, w \models \psi$
- $\mathcal{M}, w \models K\varphi$ if for each $v \in W$, if wRv, then $\mathcal{M}, v \models \varphi$

Epistemic Logic

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

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

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 $\mathcal{M}, v \models \varphi$

Results

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

Results

The logic **S5** contains the following axioms and rules:

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

Theorem

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

- Other Models
- Aumann Structures
- Group Knowledge

R. Aumann. *Interactive Epistemology I & II*. 1999.

- Let W be a set of worlds, or states.

- Let W be a set of worlds, or states.
- Let S be the set of all states of nature:

- Let W be a set of worlds, or states.
- Let S be the set of all states of nature: A state of nature is a complete description of the *exogenous* parameters (i.e. facts about the physical world) that do not depend on the agents' uncertainties.

- Let W be a set of worlds, or states.
- Let S be the set of all states of nature.
- A set $E \subseteq W$, called an **event**, is true at state w if $w \in E$.

Definition

Aumann Model An **Aumann model based on S** is a triple $\langle W, \Pi, \sigma \rangle$, where W is a nonempty set, Π is a partition over W and $\sigma : W \to S$.

Definition

Knowledge Function Let $\mathcal{M} = \langle W, \Pi, \sigma \rangle$ be an Aumann model. The **knowledge function**, $\mathsf{K} : \wp(W) \to \wp(W)$, based on $\mathcal{M}$ is defined as follows:

$$\mathsf{K}(E) = \{w \mid \Pi(w) \subseteq E\}$$

Lemma

Let $\mathcal{M} = \langle W, \Pi, \sigma \rangle$ be a Aumann model and K *the knowledge function based on* $\mathcal{M}$. *For each* $E, F \subseteq W$

$E \subseteq F \Rightarrow \mathsf{K}(E) \subseteq \mathsf{K}(F)$	Monotonicity
$\mathsf{K}(E \cap F) = \mathsf{K}(E) \cap \mathsf{K}(F)$	Closure Under Intersection
$\mathsf{K}(E) \subseteq E$	Truth
$\mathsf{K}(E) \subseteq \mathsf{K}(\mathsf{K}(E))$	Positive introspection
$\overline{\mathsf{K}(E)} \subseteq \mathsf{K}(\overline{\mathsf{K}(E)})$	Negative introspection
$\mathsf{K}(\emptyset) = \emptyset$	Consistency

where $\overline{E}$ *means the set-theoretic complement of* E *(relative to* W*).*

We can give analogous correspondence, completeness, etc. proofs.

Bayesian Structures

- Let W be a set of worlds and $\Delta(W)$ be the set of probability distributions over W.
- We are interested in functions $p : W \rightarrow \Delta(W)$.
- The basic intuition is that for each state $w \in W$, $p(w) \in \Delta(W)$ is a probability function over W.
- So, $p(w)(v)$ is the probability the agent assigns to state v in state w. To ease notation we write p_w for $p(w)$.

Definition

The pair $\langle W, p\rangle$ is called a **Bayesian frame**, where $W \neq \emptyset$ is any set, and $p : W \rightarrow \Delta(W)$ is a function such that

$$\text{if } p_w(v) > 0 \text{ then } p_w = p_v$$

Given a Bayesian frame $\mathcal{F} = \langle W, p\rangle$ and a set of states S, an **Bayesian model based on S** is a triple $\langle W, p, \sigma\rangle$, where $\sigma : W \rightarrow S$.

Definition

For each $r \in [0, 1]$ define $B^r : 2^W \rightarrow 2^W$ as follows

$$B^r(E) = \{w \mid p_w(E) \geq r\}$$

Observation: We can define a possibility model from a Bayesian model as follows. Let $\langle W, p, \sigma \rangle$ be a Bayesian model on a state space S. We define a possibility model $\langle W, P, \sigma \rangle$ base on S as follows: define $\mathcal{P} : W \rightarrow 2^W$ by

$$\mathcal{P}(w) = \{v \mid \pi_w(v) > 0\}$$

It is easy to see that $\mathcal{P}$ is serial, transitive and Euclidean.

- Other Models
- Aumann Structures
- Group Knowledge

Common Knowledge and Coordination

Suppose there are two friends Ann and Bob are on a bus separated by a crowd. Before the bus comes to the next stop a mutual friend from outside the bus yells "get off at the next stop to get a drink?".

Say Ann is standing near the front door and Bob near the back door. When the bus comes to a stop, will they get off?

D. Lewis. *Convention*. 1969.

M. Chwe. *Rational Ritual*. 2001.

Common Knowledge and Coordination

Suppose there are two friends Ann and Bob are on a bus separated by a crowd. Before the bus comes to the next stop a mutual friend from outside the bus yells "get off at the next stop to get a drink?".

Say Ann is standing near the front door and Bob near the back door. When the bus comes to a stop, will they get off?

D. Lewis. *Convention*. 1969.

M. Chwe. *Rational Ritual*. 2001.

Common Knowledge and Coordination

Suppose there are two friends Ann and Bob are on a bus separated by a crowd. Before the bus comes to the next stop a mutual friend from outside the bus yells "get off at the next stop to get a drink?".

Say Ann is standing near the front door and Bob near the back door. When the bus comes to a stop, will they get off?

D. Lewis. *Convention*. 1969.

M. Chwe. *Rational Ritual*. 2001.

Common Knowledge and Coordination

Suppose there are two friends Ann and Bob are on a bus separated by a crowd. Before the bus comes to the next stop a mutual friend from outside the bus yells "get off at the next stop to get a drink?".

Say Ann is standing near the front door and Bob near the back door. When the bus comes to a stop, will they get off?

D. Lewis. *Convention*. 1969.

M. Chwe. *Rational Ritual*. 2001.

Common Knowledge and Coordination

Suppose there are two friends Ann and Bob are on a bus separated by a crowd. Before the bus comes to the next stop a mutual friend from outside the bus yells "get off at the next stop to get a drink?".

Say Ann is standing near the front door and Bob near the back door. When the bus comes to a stop, will they get off?

D. Lewis. *Convention*. 1969.

M. Chwe. *Rational Ritual*. 2001.

Three Views of Common Knowledge

1. $\gamma :=$ i knows that φ, j knows that φ, i knows that j knows that φ, j knows that i knows that φ, i knows that j knows that i knows that φ, ...

   D. Lewis. *Convention, A Philosophical Study*. 1969.

2. $\gamma :=$ i and j know that (φ and γ)

   G. Harman. *Review of* Linguistic Behavior. Language (1977).

3. There is a *shared situation* s such that
   - s entails φ
   - s entails i knows φ
   - s entails j knows φ

   H. Clark and C. Marshall. *Definite Reference and Mutual Knowledge*. 1981.

J. Barwise. *Three views of Common Knowledge*. TARK (1987).

Three Views of Common Knowledge

1. $\gamma :=$ i knows that φ, j knows that φ, i knows that j knows that φ, j knows that i knows that φ, i knows that j knows that i knows that φ, ...

   D. Lewis. *Convention, A Philosophical Study*. 1969.

2. $\gamma :=$ i and j know that (φ and γ)

   G. Harman. *Review of* Linguistic Behavior. Language (1977).

3. There is a *shared situation* s such that
   - s entails φ
   - s entails i knows φ
   - s entails j knows φ

   H. Clark and C. Marshall. *Definite Reference and Mutual Knowledge*. 1981.

J. Barwise. *Three views of Common Knowledge*. TARK (1987).

Three Views of Common Knowledge

1. $\gamma :=$ *i* knows that φ, *j* knows that φ, *i* knows that *j* knows that φ, *j* knows that *i* knows that φ, *i* knows that *j* knows that *i* knows that φ, ...

   D. Lewis. *Convention, A Philosophical Study*. 1969.

2. $\gamma :=$ *i* and *j* know that (φ and γ)

   G. Harman. *Review of* Linguistic Behavior. Language (1977).

3. There is a *shared situation s* such that
   - *s* entails φ
   - *s* entails *i* knows φ
   - *s* entails *j* knows φ

   H. Clark and C. Marshall. *Definite Reference and Mutual Knowledge*. 1981.

J. Barwise. *Three views of Common Knowledge*. TARK (1987).

Common Knowledge: Iterated View

Recall $K_i : \wp(W) \rightarrow \wp(W)$ is i's knowledge function.

Define $K^m : \wp(W) \rightarrow \wp(W)$ for $m \geq 1$ by

- $K^1 E \;:=\; \bigcap_{i \in \mathcal{A}} K_i E$
- $K^{m+1} E \;:=\; K^1(K^m(E))$

$K^1 E$ means *everyone knows E*
$K^2 E$ means *everyone knows that everyone knows E*

Define $K^\infty : \wp(W) \rightarrow \wp(W)$

$$K^\infty E \;:=\; K^1 E \cap K^2 E \cap \cdots \cap K^m E \cap \cdots$$

Common Knowledge: Iterated View

Recall $K_i : \wp(W) \to \wp(W)$ is i's knowledge function.

Define $K^m : \wp(W) \to \wp(W)$ for $m \geq 1$ by

- $K^1 E \; := \; \bigcap_{i \in \mathcal{A}} K_i E$
- $K^{m+1} E \; := \; K^1(K^m(E))$

$K^1 E$ means *everyone knows E*
$K^2 E$ means *everyone knows that everyone knows E*

Define $K^\infty : \wp(W) \to \wp(W)$

$$K^\infty E \; := \; K^1 E \cap K^2 E \cap \cdots \cap K^m E \cap \cdots$$

Common Knowledge: Iterated View

Recall $K_i : \wp(W) \rightarrow \wp(W)$ is i's knowledge function.

Define $K^m : \wp(W) \rightarrow \wp(W)$ for $m \geq 1$ by

- $K^1 E \ := \ \bigcap_{i \in \mathcal{A}} K_i E$
- $K^{m+1} E \ := \ K^1(K^m(E))$

$K^1 E$ means *everyone knows E*
$K^2 E$ means *everyone knows that everyone knows E*

Define $K^\infty : \wp(W) \rightarrow \wp(W)$

$$K^\infty E \ := \ K^1 E \cap K^2 E \cap \cdots \cap K^m E \cap \cdots$$

Common Knowledge: Iterated View

Recall $K_i : \wp(W) \rightarrow \wp(W)$ is i's knowledge function.

Define $K^m : \wp(W) \rightarrow \wp(W)$ for $m \geq 1$ by

- $K^1 E \ := \ \bigcap_{i \in \mathcal{A}} K_i E$
- $K^{m+1} E \ := \ K^1(K^m(E))$

$K^1 E$ means *everyone knows* E
$K^2 E$ means *everyone knows that everyone knows* E

Define $K^\infty : \wp(W) \rightarrow \wp(W)$

$$K^\infty E \ := \ K^1 E \cap K^2 E \cap \cdots \cap K^m E \cap \cdots$$

Common Knowledge: Iterated View

Recall $K_i : \wp(W) \to \wp(W)$ is i's knowledge function.

Define $K^m : \wp(W) \to \wp(W)$ for $m \geq 1$ by

- $K^1 E \;:=\; \bigcap_{i \in \mathcal{A}} K_i E$
- $K^{m+1} E \;:=\; K^1(K^m(E))$

$K^1 E$ means *everyone knows E*
$K^2 E$ means *everyone knows that everyone knows E*

Define $K^\infty : \wp(W) \to \wp(W)$

$$K^\infty E \;:=\; K^1 E \cap K^2 E \cap \cdots \cap K^m E \cap \cdots$$

Fact Prove that for all $i \in \mathcal{A}$ and $E \subseteq W$, $\mathsf{K}_i\mathsf{K}^\infty(E) = \mathsf{K}^\infty(E)$.

Fact Prove that for all $i \in \mathcal{A}$ and $E \subseteq W$, $\mathsf{K}_i\mathsf{K}^\infty(E) = \mathsf{K}^\infty(E)$.

Suppose you are told "Ann and Bob are going together,"' and respond "sure, that's common knowledge." What you mean is not only that everyone knows this, but also that the announcement is pointless, occasions no surprise, reveals nothing new; in effect, that the situation after the announcement does not differ from that before. ...the event "Ann and Bob are going together" — call it E — is common knowledge if and only if some event — call it F — happened that entails E and also entails all players' knowing F (like all players met Ann and Bob at an intimate party). *(Aumann, pg. 271, footnote 8)*

Definition

Self-Evident Event An event F is **self-evident** if $\mathsf{K}_i(F) = F$ for all $i \in \mathcal{A}$.

Definition

Knowledge Field Let $\langle W, \{\Pi_i\}_{i\in\mathcal{A}}, \sigma\rangle$ be a multi-agent Aumann model. For each $i \in \mathcal{A}$, the **knowledge field of** i, denoted $\mathbb{K}_i$, is the family of all unions of cells in Π_i.

Lemma

An event E is commonly known iff some self-evident event that entails E obtains. Formally, $\mathsf{K}^\infty(E)$ is the largest event in $\cap_{i\in\mathcal{A}}\mathbb{K}_i$ that is included in E.

Common Knowledge in Epistemic Logic

Definition

The operator "everyone knows φ", denoted $E\varphi$, is defined as follows

$$E\varphi \quad := \quad \bigwedge_{i\in\mathcal{A}} K_i\varphi$$

Common Knowledge in Epistemic Logic

Definition

The operator "everyone knows φ", denoted $E\varphi$, is defined as follows

$$E\varphi \quad := \quad \bigwedge_{i \in \mathcal{A}} K_i\varphi$$

Definition

The multi-agent epistemic language with common knowledge is generated by the following grammar:

$$p \mid \neg\varphi \mid \varphi \wedge \varphi \mid K_i\varphi \mid C\varphi$$

where $p \in \mathsf{At}$ and $i \in \mathcal{A}$.

Common Knowledge in Epistemic Logic

Definition

The truth of $C\varphi$ is:

$$\mathcal{M}, w \models C\varphi \quad \text{iff} \quad \text{for all } v \in W, \text{ if } wR^*v \text{ then } \mathcal{M}, v \models \varphi$$

where $R^* := (\bigcup_{i \in \mathcal{A}} R_i)^*$ is the reflexive transitive closure of the union of the R_i's.

Common Knowledge in Epistemic Logic

Definition

The truth of $C\varphi$ is:

$$\mathcal{M}, w \models C\varphi \quad \text{iff} \quad \text{for all } v \in W, \text{ if } wR^*v \text{ then } \mathcal{M}, v \models \varphi$$

where $R^* := (\bigcup_{i \in \mathcal{A}} R_i)^*$ is the reflexive transitive closure of the union of the R_i's.

$\mathcal{M}, w \models C\varphi$ iff every finite path starting at w ends with a state satisfying φ.

Example

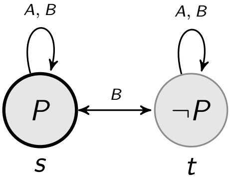

P means "The talk is at 2PM".

Example

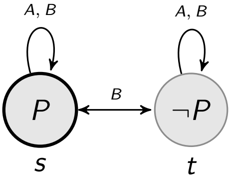

P means "The talk is at 2PM".

$\mathcal{M}, s \models \neg K_B P$

Example

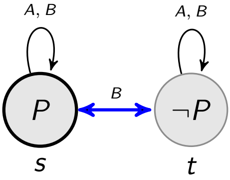

P means "The talk is at 2PM".

$\mathcal{M}, s \models \neg K_B P$

Example

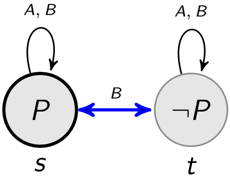

P means "The talk is at 2PM".

$\mathcal{M}, s \models E\neg K_B P$

Example

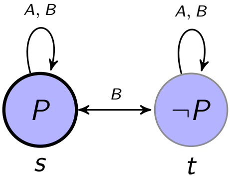

P means "The talk is at 2PM".

$\mathcal{M}, s \models E\neg K_B P$

Example

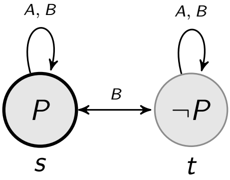

P means "The talk is at 2PM".

$\mathcal{M}, s \models C\neg K_B P$

Another Example

Two players Ann and Bob are told that the following will happen. Some positive integer n will be chosen and *one* of n, $n+1$ will be written on Ann's forehead, the other on Bob's. Each will be able to see the other's forehead, but not his/her own.

Suppose the number are (2,3).

Do the agents know there numbers are less than 1000?

Is it common knowledge that their numbers are less than 1000?

Another Example

Two players Ann and Bob are told that the following will happen. Some positive integer n will be chosen and *one* of n, $n+1$ will be written on Ann's forehead, the other on Bob's. Each will be able to see the other's forehead, but not his/her own.

Suppose the number are (2,3).

Do the agents know there numbers are less than 1000?

Is it common knowledge that their numbers are less than 1000?

Another Example

Two players Ann and Bob are told that the following will happen. Some positive integer n will be chosen and *one* of n, $n+1$ will be written on Ann's forehead, the other on Bob's. Each will be able to see the other's forehead, but not his/her own.

Suppose the number are (2,3).

Do the agents know there numbers are less than 1000?

Is it common knowledge that their numbers are less than 1000?

Another Example

Two players Ann and Bob are told that the following will happen. Some positive integer n will be chosen and *one* of n, $n+1$ will be written on Ann's forehead, the other on Bob's. Each will be able to see the other's forehead, but not his/her own.

Suppose the number are (2,3).

Do the agents know there numbers are less than 1000?

Is it common knowledge that their numbers are less than 1000?

Logics with Common Knowledge

The following axiom and rule need to be added to **S5** to deal with the common knowledge operator:

- $C(\varphi \rightarrow \psi) \rightarrow (C\varphi \rightarrow C\psi)$
- $C\varphi \rightarrow (\varphi \wedge EC\varphi)$
- $C(\varphi \rightarrow E\varphi) \rightarrow (\varphi \rightarrow C\varphi)$

Theorem

$\mathbf{S5}^C$ *is sound and weakly complete with respect to the class of all Kripke frames where the relations are equivalence relations.*

Logics with Common Knowledge

The following axiom and rule need to be added to **S5** to deal with the common knowledge operator:

- $C(\varphi \to \psi) \to (C\varphi \to C\psi)$
- $C\varphi \to (\varphi \wedge EC\varphi)$
- $C(\varphi \to E\varphi) \to (\varphi \to C\varphi)$

Theorem

$\mathbf{S5}^C$ *is sound and weakly complete with respect to the class of all Kripke frames where the relations are equivalence relations.*

Concluding Remark

- In Relational Models (including Aumann Structures), the first two views of Common Knowledge are mathematically equivalent and it is not clear how to represent the third view.
- They can be separated using alternative semantics.

  J. van Benthem and D. Sarenac. *The Geometry of Knowledge*. 2005.

  J. Barwise. *Three Views of Common Knowledge*. TARK, 1987.
- There are other notions of "group knowledge": distributed knowledge
- In many social situations, other levels of knowledge are of interest...

Concluding Remark

- In Relational Models (including Aumann Structures), the first two views of Common Knowledge are mathematically equivalent and it is not clear how to represent the third view.
- They can be separated using alternative semantics.

  J. van Benthem and D. Sarenac. *The Geometry of Knowledge*. 2005.

  J. Barwise. *Three Views of Common Knowledge*. TARK, 1987.
- There are other notions of "group knowledge": distributed knowledge
- In many social situations, other levels of knowledge are of interest...

Concluding Remark

- In Relational Models (including Aumann Structures), the first two views of Common Knowledge are mathematically equivalent and it is not clear how to represent the third view.

- They can be separated using alternative semantics.

  J. van Benthem and D. Sarenac. *The Geometry of Knowledge*. 2005.

  J. Barwise. *Three Views of Common Knowledge*. TARK, 1987.

- There are other notions of "group knowledge": distributed knowledge

- In many social situations, other levels of knowledge are of interest...

Concluding Remark

- In Relational Models (including Aumann Structures), the first two views of Common Knowledge are mathematically equivalent and it is not clear how to represent the third view.
- They can be separated using alternative semantics.

  J. van Benthem and D. Sarenac. *The Geometry of Knowledge*. 2005.

  J. Barwise. *Three Views of Common Knowledge*. TARK, 1987.
- There are other notions of "group knowledge": distributed knowledge
- In many social situations, other levels of knowledge are of interest...

Concluding Remark

- In Relational Models (including Aumann Structures), the first two views of Common Knowledge are mathematically equivalent and it is not clear how to represent the third view.

- They can be separated using alternative semantics.

  J. van Benthem and D. Sarenac. *The Geometry of Knowledge*. 2005.

  J. Barwise. *Three Views of Common Knowledge*. TARK, 1987.

- There are other notions of "group knowledge": distributed knowledge

- In many social situations, other levels of knowledge are of interest...

Time for a beak.