Logics of Rational Agency

Lecture 2

Eric Pacuit

Tilburg Institute for Logic and Philosophy of Science
Tilburg Univeristy
ai.stanford.edu/~epacuit

June 23, 2010

Part 2: Ingredients of a Logical Analysis of Rational Agency

- Logics of Informational Attitudes and Informative Actions
- Logics of Motivational Attitudes (Preferences)
- Time, Action and Agency

Basic Ingredients

- What are the basic building blocks? (the nature of time (continuous or discrete/branching or linear), how (primitive) *events* or *actions* are represented, how *causal* relationships are represented and what constitutes a *state of affairs*.)

Basic Ingredients

- What are the basic building blocks? (the nature of time (continuous or discrete/branching or linear), how (primitive) *events* or *actions* are represented, how *causal* relationships are represented and what constitutes a *state of affairs*.)
- Single agent vs. many agents.

Basic Ingredients

- What are the basic building blocks? (the nature of time (continuous or discrete/branching or linear), how (primitive) *events* or *actions* are represented, how *causal* relationships are represented and what constitutes a *state of affairs*.)
- Single agent vs. many agents.
- What the the primitive operators?
  - Informational attitudes
  - Motivational attitudes
  - Normative attitudes

Basic Ingredients

- What are the basic building blocks? (the nature of time (continuous or discrete/branching or linear), how (primitive) *events* or *actions* are represented, how *causal* relationships are represented and what constitutes a *state of affairs*.)
- Single agent vs. many agents.
- What the the primitive operators?
  - Informational attitudes
  - Motivational attitudes
  - Normative attitudes
- Static vs. dynamic

✓ informational attitudes (eg., knowledge, belief, certainty)

✓ group notions (eg., common knowledge and coalitional ability)

✓ time, actions and ability

✓ motivational attitudes (eg., preferences)

✓ normative attitudes (eg., obligations)

Logics of Informational Attitudes and Informative Actions

See the courses:

1. Tutorial on Epistemic and Modal Logic by Hans van Ditmarsch
2. Dynamic Epistemic Logic by Hans van Ditmarsch
3. Multi-Agent Belief Dynamics by Alexandru Baltag and Sonja Smets

Also,

EP. *Logics of Informational Attitudes and Informative Actions*. to appear, 2010 (available on the course website).

Logics of Informational Attitudes and Informative Actions

See the courses:

1. Tutorial on Epistemic and Modal Logic by Hans van Ditmarsch
2. Dynamic Epistemic Logic by Hans van Ditmarsch
3. Multi-Agent Belief Dynamics by Alexandru Baltag and Sonja Smets

Also,

EP. *Logics of Informational Attitudes and Informative Actions*. to appear, 2010 (available on the course website).

Rather than a general introduction, we present results not typically discussed in introductions to epistemic logic:

1. Can we agree to disagree?
2. How many levels of knowledge are there?

Agreeing to Disagree

Theorem: Suppose that n agents share a common prior and have different private information. If there is common knowledge in the group of the posterior probabilities, then the posteriors must be equal.

Robert Aumann. *Agreeing to Disagree*. Annals of Statistics **4** (1976).

Agreeing to Disagree

Theorem: Suppose that *n* agents share a common prior and have different private information. If there is common knowledge in the group of the posterior probabilities, then the posteriors must be equal.

Robert Aumann. *Agreeing to Disagree*. Annals of Statistics **4** (1976).

G. Bonanno and K. Nehring. *Agreeing to Disagree: A Survey*. (manuscript) 1997.

2 Scientists Perform an Experiment

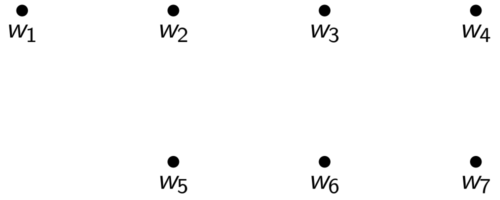

They agree the true state is one of seven different states.

2 Scientists Perform an Experiment

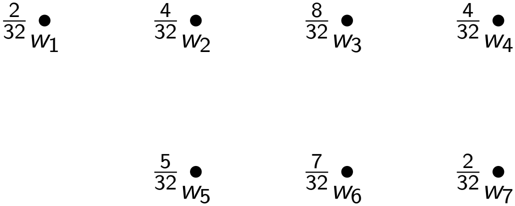

They agree on a common prior.

2 Scientists Perform an Experiment

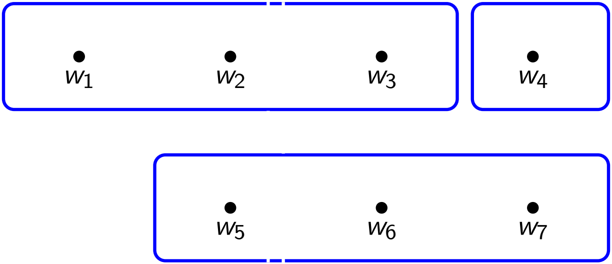

They agree that Experiment 1 would produce the blue partition.

2 Scientists Perform an Experiment

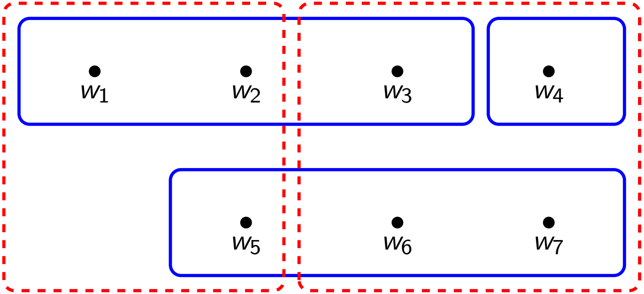

They agree that Experiment 1 would produce the blue partition and Experiment 2 the red partition.

2 Scientists Perform an Experiment

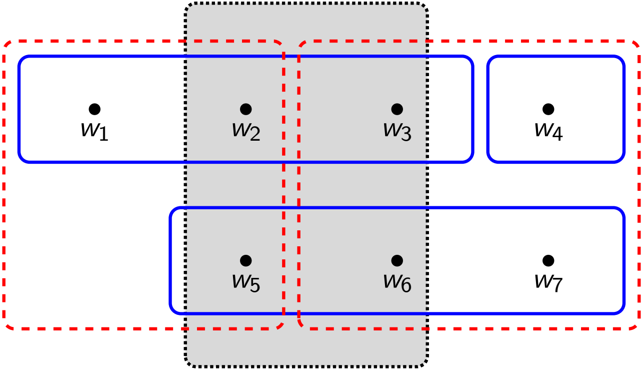

They are interested in the truth of $E = \{w_2, w_3, w_5, w_6\}$.

2 Scientists Perform an Experiment

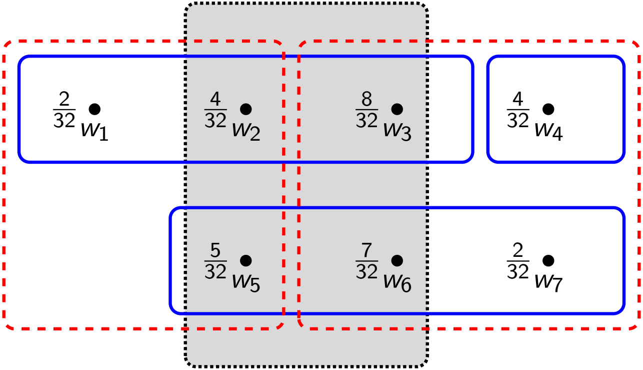

So, they agree that $P(E) = \frac{24}{32}$.

2 Scientists Perform an Experiment

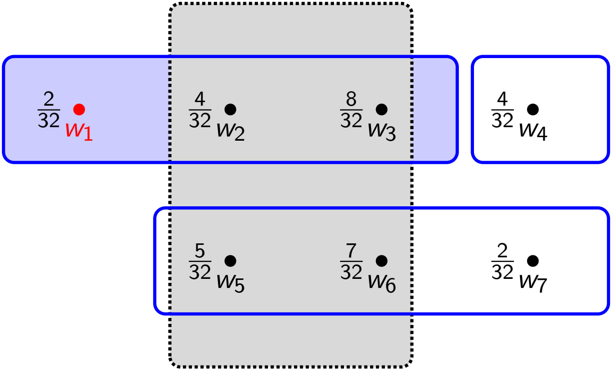

Also, that *if the true state is* w_1, then Experiment 1 will yield

$$P(E|I) = \frac{P(E \cap I)}{P(I)} = \frac{12}{14}$$

2 Scientists Perform an Experiment

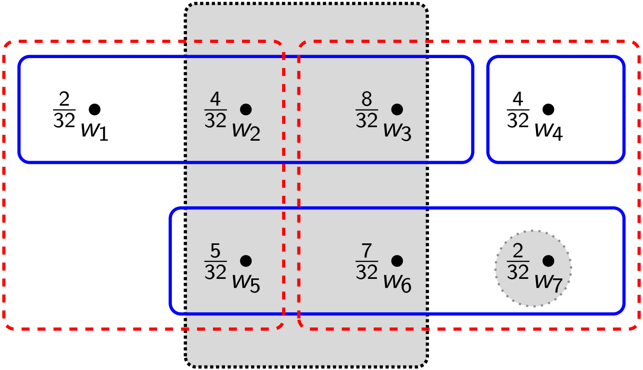

Suppose the true state is w_7 and the agents preform the experiments.

2 Scientists Perform an Experiment

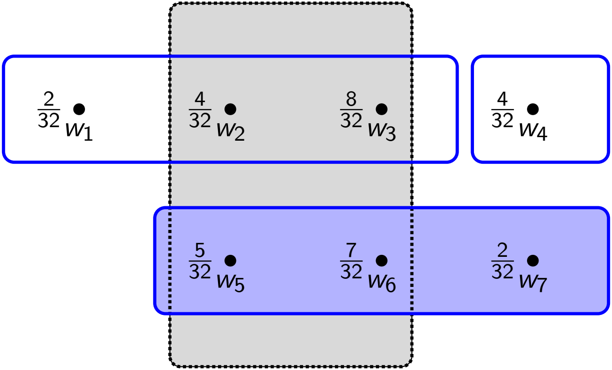

Suppose the true state is w_7, then $Pr_1(E) = \frac{12}{14}$

2 Scientists Perform an Experiment

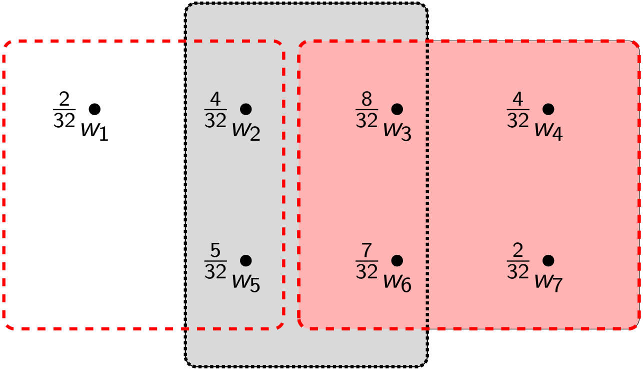

Then $Pr_1(E) = \frac{12}{14}$ and $Pr_2(E) = \frac{15}{21}$

2 Scientists Perform an Experiment

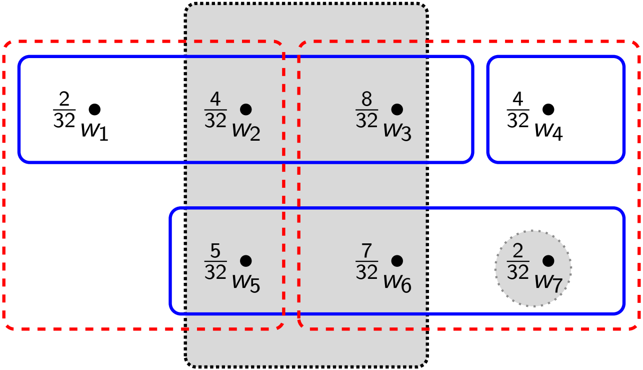

Suppose they exchange emails with the new subjective probabilities: $Pr_1(E) = \frac{12}{14}$ and $Pr_2(E) = \frac{15}{21}$

2 Scientists Perform an Experiment

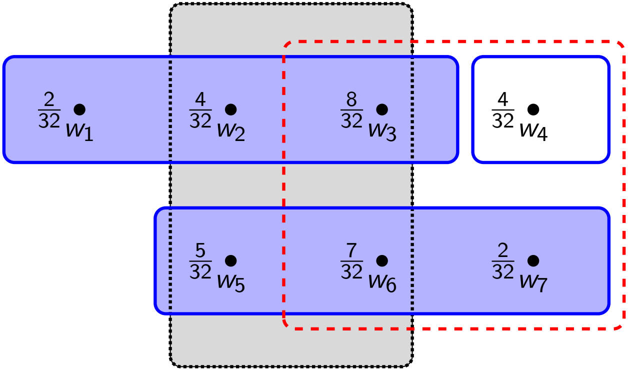

Agent 2 learns that w_4 is **NOT** the true state (same for Agent 1).

2 Scientists Perform an Experiment

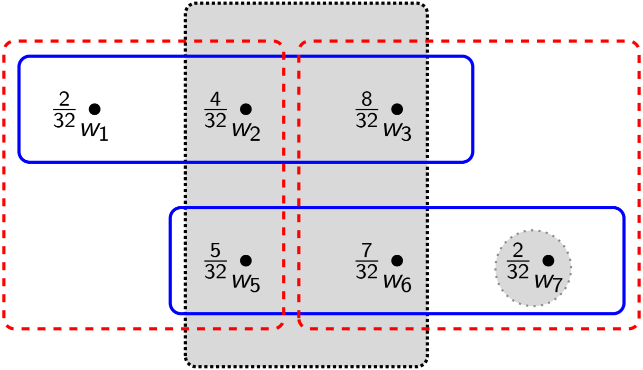

Agent 2 learns that w_4 is **NOT** the true state (same for Agent 1).

2 Scientists Perform an Experiment

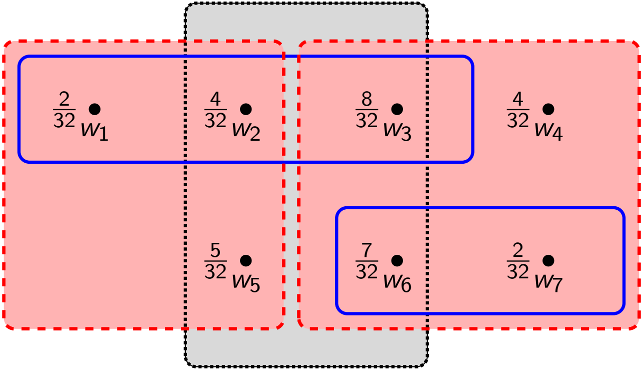

Agent 1 learns that w_5 is **NOT** the true state (same for Agent 1).

2 Scientists Perform an Experiment

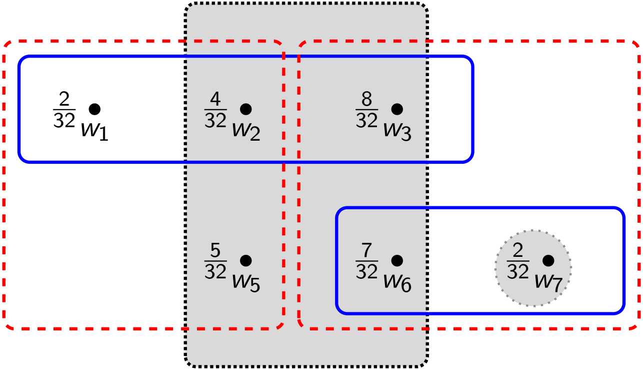

The new probabilities are $Pr_1(E|I') = \frac{7}{9}$ and $Pr_2(E|I') = \frac{15}{17}$

2 Scientists Perform an Experiment

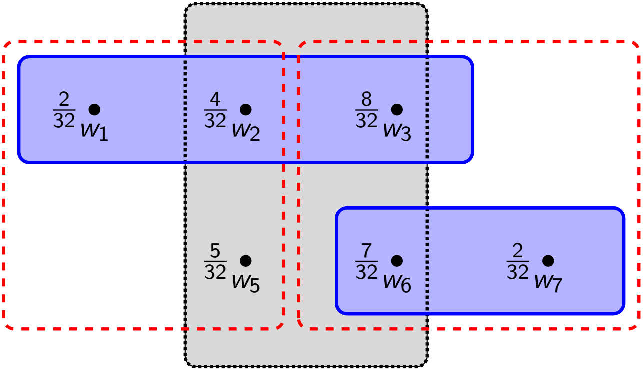

After exchanging this information ($Pr_1(E|I') = \frac{7}{9}$ and $Pr_2(E|I') = \frac{15}{17}$), Agent 2 learns that w_3 is **NOT** the true state.

2 Scientists Perform an Experiment

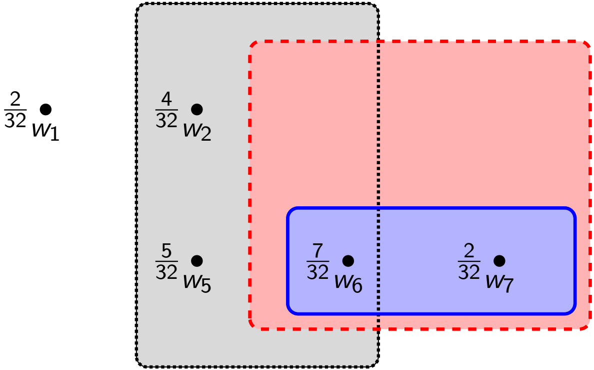

No more revisions are possible and the agents agree on the posterior probabilities.

Formal Models of Knowledge

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

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.

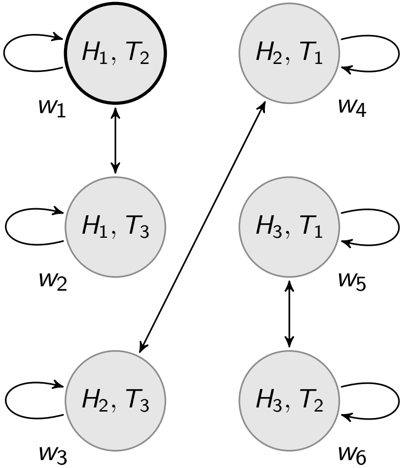

Example

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.

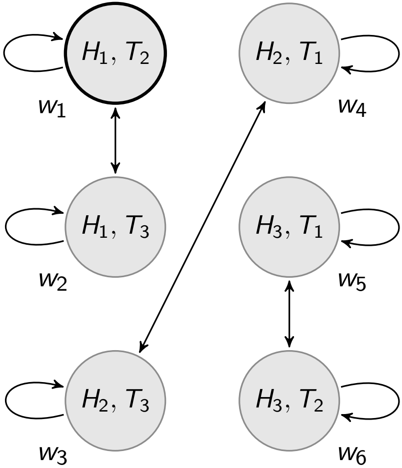

Example

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.

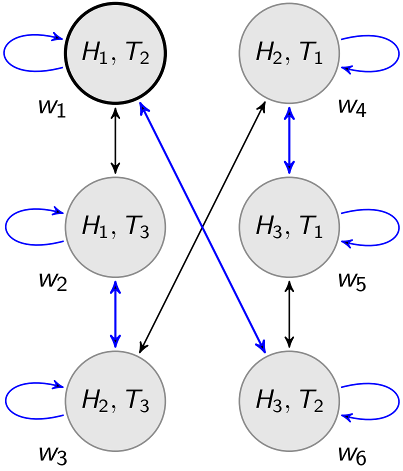

Example

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.

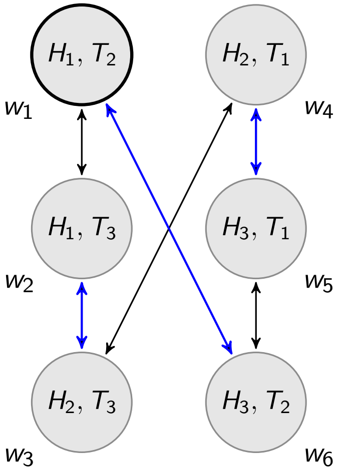

Example

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.

$\mathcal{M}, w \models K_B(K_A H_1 \vee K_A \neg H_1)$

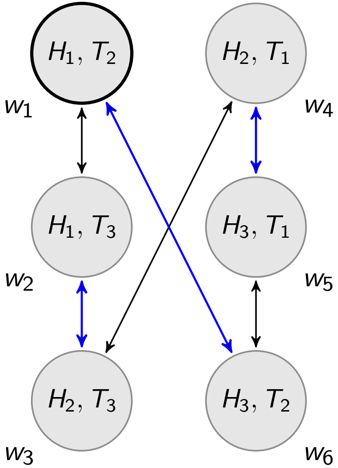

Example

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.

$\mathcal{M}, w \models K_B(K_A H_1 \vee K_A \neg H_1)$

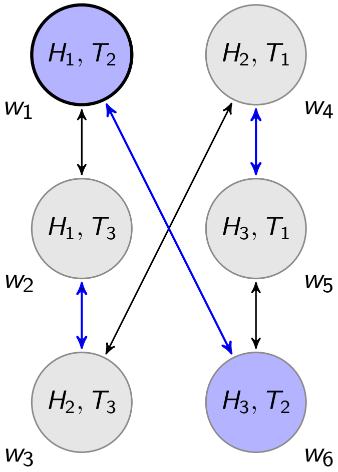

Example

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.

$\mathcal{M}, w \models K_B(K_A H_1 \vee K_A \neg H_1)$

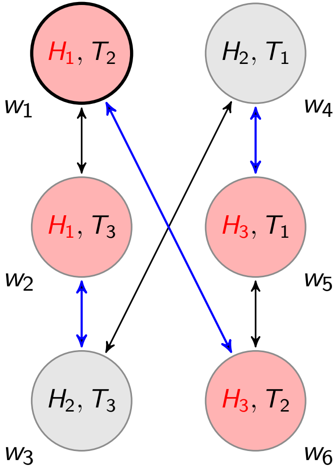

Group Knowledge

$K_A P$: "Ann knows that P"

Group Knowledge

$K_A P$: "Ann knows that P"

$K_B P$: "Bob knows that P"

Group Knowledge

$K_A P$: "Ann knows that P"

$K_B P$: "Bob knows that P"

$K_A K_B P$: "Ann knows that Bob knows that P"

Group Knowledge

$K_A P$: "Ann knows that P"

$K_B P$: "Bob knows that P"

$K_A K_B P$: "Ann knows that Bob knows that P"

$K_A P \wedge K_B P$: "Every one knows P".

Group Knowledge

$K_A P$: "Ann knows that P"

$K_B P$: "Bob knows that P"

$K_A K_B P$: "Ann knows that Bob knows that P"

$K_A P \wedge K_B P$: "Every one knows P". let $EP := K_A P \wedge K_B P$

Group Knowledge

$K_A P$: "Ann knows that P"

$K_B P$: "Bob knows that P"

$K_A K_B P$: "Ann knows that Bob knows that P"

$K_A P \wedge K_B P$: "Every one knows P". let $EP := K_A P \wedge K_B P$

$K_A EP$: "Ann knows that everyone knows that P".

Group Knowledge

$K_A P$: "Ann knows that P"

$K_B P$: "Bob knows that P"

$K_A K_B P$: "Ann knows that Bob knows that P"

$K_A P \wedge K_B P$: "Every one knows P". let $EP := K_A P \wedge K_B P$

$K_A EP$: "Ann knows that everyone knows that P".

EEP: "Everyone knows that everyone knows that P".

Group Knowledge

$K_A P$: "Ann knows that P"

$K_B P$: "Bob knows that P"

$K_A K_B P$: "Ann knows that Bob knows that P"

$K_A P \wedge K_B P$: "Every one knows P". let $EP := K_A P \wedge K_B P$

$K_A EP$: "Ann knows that everyone knows that P".

EEP: "Everyone knows that everyone knows that P".

$EEEP$: "Everyone knows that everyone knows that everyone knows that P."

Group Knowledge

$K_A P$: "Ann knows that P"

$K_B P$: "Bob knows that P"

$K_A K_B P$: "Ann knows that Bob knows that P"

$K_A P \wedge K_B P$: "Every one knows P". let $EP := K_A P \wedge K_B P$

$K_A EP$: "Ann knows that everyone knows that P".

EEP: "Everyone knows that everyone knows that P".

$EEEP$: "Everyone knows that everyone knows that everyone knows that P."

CP: "it is common knowledge that P"

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, $\ldots$

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

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).

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).

Dissecting Aumann's Theorem

- "No Trade" Theorems (Milgrom and Stokey); from probabilities of events to aggregates (McKelvey and Page); Common Prior Assumption, etc.

Dissecting Aumann's Theorem

- "No Trade" Theorems (Milgrom and Stokey); from probabilities of events to aggregates (McKelvey and Page); Common Prior Assumption, etc.

- How do the posteriors *become* common knowledge?

J. Geanakoplos and H. Polemarchakis. *We Can't Disagree Forever*. Journal of Economic Theory (1982).

Geanakoplos and Polemarchakis

Revision Process: Given event A, 1: "My probability of A is q", 2: "My probability of A is r, 1: "My probability of A is now q', 2: "My probability of A is now r', etc.

Geanakoplos and Polemarchakis

Revision Process: Given event A, 1: "My probability of A is q", 2: "My probability of A is r, 1: "My probability of A is now q', 2: "My probability of A is now r', etc.

- Assuming that the information partitions are finite, given an event A, the revision process converges in finitely many steps.

Geanakoplos and Polemarchakis

Revision Process: Given event A, 1: "My probability of A is q", 2: "My probability of A is r, 1: "My probability of A is now q', 2: "My probability of A is now r', etc.

- Assuming that the information partitions are finite, given an event A, the revision process converges in finitely many steps.
- For each n, there are examples where the process takes n steps.

Geanakoplos and Polemarchakis

Revision Process: Given event A, 1: "My probability of A is q", 2: "My probability of A is r, 1: "My probability of A is now q', 2: "My probability of A is now r', etc.

- Assuming that the information partitions are finite, given an event A, the revision process converges in finitely many steps.
- For each n, there are examples where the process takes n steps.
- An *indirect communication* equilibrium is not necessarily a *direct communication* equilibrium.

Parikh and Krasucki

Protocol: graph on the set of agents specifying the legal pairwise communication (who can talk to who).

Parikh and Krasucki

Protocol: graph on the set of agents specifying the legal pairwise communication (who can talk to who).

- If the protocol is **fair**, then the *limiting probability* of an event A will be the same for all agents in the group.

Parikh and Krasucki

Protocol: graph on the set of agents specifying the legal pairwise communication (who can talk to who).

- If the protocol is **fair**, then the *limiting probability* of an event A will be the same for all agents in the group.
- Consensus can be reached without common knowledge: "everyone must know the common prices of commodities; however, it does not make sense to demand that everyone knows the details of every commercial transaction."

Qualitative Generalizations

Assuming a version of *Savage's Sure-Thing Principle*, their cannot be common knowledge that two-like minded individuals make different decisions.

The Sure-Thing Principle

A businessman contemplates buying a certain piece of property. He considers the outcome of the next presidential election relevant.

The Sure-Thing Principle

> *A businessman contemplates buying a certain piece of property. He considers the outcome of the next presidential election relevant. So, to clarify the matter to himself, he asks whether he would buy if he* knew *that the Democratic candidate were going to win, and decides that he would.*

The Sure-Thing Principle

A businessman contemplates buying a certain piece of property. He considers the outcome of the next presidential election relevant. So, to clarify the matter to himself, he asks whether he would buy if he knew *that the Democratic candidate were going to win, and decides that he would. Similarly, he considers whether he would buy if he* knew *a Republican candidate were going to win, and again he finds that he would.*

The Sure-Thing Principle

> *A businessman contemplates buying a certain piece of property. He considers the outcome of the next presidential election relevant. So, to clarify the matter to himself, he asks whether he would buy if he* knew *that the Democratic candidate were going to win, and decides that he would. Similarly, he considers whether he would buy if he* knew *a Republican candidate were going to win, and again he finds that he would. Seeing that he would buy in either event, he decides that he should buy, even though he does not know which event obtains, or will obtain, as we would ordinarily say.* (Savage, 1954)

Dissecting Aumann's Theorem

- Qualitative versions: like-minded individuals cannot agree to make different decisions.

M. Bacharach. *Some Extensions of a Claim of Aumann in an Axiomatic Model of Knowledge*. Journal of Economic Theory (1985).

J.A.K. Cave. *Learning to Agree*. Economic Letters (1983).

D. Samet. *The Sure-Thing Principle and Independence of Irrelevant Knowledge*. 2008.

Analyzing Agreement Theorems in Dynamic Epistemic/Doxastic Logic

C. Degremont and O. Roy. *Agreement Theorems in Dynamic-Epistemic Logic*. in A. Heifetz (ed.), Proceedings of TARK XI, 2009, pp.91-98.

L. Demey. *Agreeing to Disagree in Probabilistic Dynamic Epistemic Logic*. ILLC, Masters Thesis, forthcoming.

Levels of Knowledge

Levels of Knowledge

What are the *states of knowledge* created in a group when communication takes place? What happens when communication is not the the whole group, but pairwise?

R. Parikh and P. Krasucki. *Communication, Consensus and Knowledge*. Journal of Economic Theory (1990).

Informal Definition: Given some fact P and a set of agents $\mathcal{A}$, a **state of knowledge** is a (consistent) description of the agents first-order and higher-order information about P.

Levels of Knowledge

What are the *states of knowledge* created in a group when communication takes place? What happens when communication is not the the whole group, but pairwise?

R. Parikh and P. Krasucki. *Communication, Consensus and Knowledge*. Journal of Economic Theory (1990).

Informal Definition: Given some fact P and a set of agents $\mathcal{A}$, a **state of knowledge** is a (consistent) description of the agents first-order and higher-order information about P.

For example: $\{P, K_A P, K_B P, K_B K_A P\}$

Levels of Knowledge

Informal Definition: Given some fact P and a set of agents $\mathcal{A}$, a **state of knowledge** is a (consistent) description of the agents first-order and higher-order information about P.

For example: $\{P, K_A P, K_B P, K_B K_A P\}$

Also called **level of knowledge**, **hierarchy of knowledge**.

Levels of Knowledge

Informal Definition: Given some fact P and a set of agents $\mathcal{A}$, a **state of knowledge** is a (consistent) description of the agents first-order and higher-order information about P.

For example: $\{P, K_A P, K_B P, K_B K_A P\}$

Also called **level of knowledge**, **hierarchy of knowledge**.

At one extreme, no one may have any information about P and the other extreme is when there is common knowledge of P.

There are many interesting levels in between...

Levels of Knowledge

Informal Definition: Given some fact P and a set of agents $\mathcal{A}$, a **state of knowledge** is a (consistent) description of the agents first-order and higher-order information about P.

For example: $\{P, K_iP, K_jP, K_jK_iP\}$

Also called **level of knowledge**, **hierarchy of knowledge**.

At one extreme, no one may have any information about P and the other extreme is when there is common knowledge of P.

There are many interesting levels in between...

Levels of Knowledge

Informal Definition: Given some fact P and a set of agents $\mathcal{A}$, a **state of belief** is a (consistent) description of the agents first-order and higher-order information about P.

For example: $\{P, B_iP, B_jP, B_jB_iP\}$

Also called **level of belief**, **hierarchy of belief**.

At one extreme, no one may have any information about P and the other extreme is when there is common belief of P.

There are many interesting levels in between...

Some Questions/Issues

- How do states of knowledge influence decisions in *game situations*?
- Can we *realize* any state of knowledge?
- What is a *state* in an epistemic model?
- Is an epistemic model *common knowledge* among the agents?

States of Knowledge in Games

R. Parikh. *Levels of knowledge, games and group action*. Research in Economics 57, pp. 267 - 281 (2003).

States of Knowledge in Games

	G	N
g		
n		

States of Knowledge in Games

	G	N
g		$(1,0)$
n	$(0,1)$	

States of Knowledge in Games

	G	N
g		$(1,0)$
n	$(0,1)$	$(0,0)$

States of Knowledge in Games

	G	N
g	$(-100, -10)$	$(1, 0)$
n	$(0, 1)$	$(0, 0)$

States of Knowledge in Games

	G	N
g	$(-100, -10)$	$(1, 0)$
n	$(0, 1)$	$(0, 0)$

States of Knowledge in Games

	G	N
g	$(-100, -10)$	$(1, 0)$
n	$(0, 1)$	$(0, 0)$

States of Knowledge in Games

	G	N
g	$(-100, -10)$	$(1, 0)$
n	$(0, 1)$	$(0, 0)$

$$C_{p,m}c$$

States of Knowledge in Games

	G	N
g	$(-100, -10)$	$(1, 0)$
n	$(0, 1)$	$(0, 0)$

$$K_p c,\ \neg K_m K_p c$$

Realizing States of Knowledge

What is the exact relationship between

Realizing States of Knowledge

What is the exact relationship between 1. types of communicatory events (public broadcasts, private announcements, DEL events, etc.),

Realizing States of Knowledge

What is the exact relationship between 1. types of communicatory events (public broadcasts, private announcements, DEL events, etc.), 2. the *protocol* (eg., synchronous, asynchronous, communication graph) and

Realizing States of Knowledge

What is the exact relationship between 1. types of communicatory events (public broadcasts, private announcements, DEL events, etc.), 2. the *protocol* (eg., synchronous, asynchronous, communication graph) and 3. different levels of knowledge?

Realizing States of Knowledge

What is the exact relationship between 1. types of communicatory events (public broadcasts, private announcements, DEL events, etc.), 2. the *protocol* (eg., synchronous, asynchronous, communication graph) and 3. different levels of knowledge?

Classic example: general's problem or the email problem show that common knowledge cannot be realized in systems with only asynchronous communication.
(cf. Halpern and Moses, Rubinstein)

Realizing States of Knowledge

What is the exact relationship between 1. types of communicatory events (public broadcasts, private announcements, DEL events, etc.), 2. the *protocol* (eg., synchronous, asynchronous, communication graph) and 3. different levels of knowledge?

Classic example: general's problem or the email problem show that common knowledge cannot be realized in systems with only asynchronous communication.
(cf. Halpern and Moses, Rubinstein)

What about other levels of knowledge?

R. Parikh and P. Krasucki. *Levels of knowledge in distributed computing*. Sadhana-Proceedings of the Indian Academy of Science 17 (1992).

What is a *State*?

Possible worlds, or states, are taken as primitive in Kripke structures. But in many applications, we intuitively understand what a state *is*:

What is a *State*?

Possible worlds, or states, are taken as primitive in Kripke structures. But in many applications, we intuitively understand what a state *is*:

Dynamic logic: a program state (assignment of values to variables)
Temporal logic: a moment in time
Distributed system: a sequence of local states for each process

What is a *State*?

Possible worlds, or states, are taken as primitive in Kripke structures. But in many applications, we intuitively understand what a state *is*:

Dynamic logic: a program state (assignment of values to variables)
Temporal logic: a moment in time
Distributed system: a sequence of local states for each process

What about in *game situations*?

What is a *State*?

Possible worlds, or states, are taken as primitive in Kripke structures. But in many applications, we intuitively understand what a state *is*:

Dynamic logic: a program state (assignment of values to variables)
Temporal logic: a moment in time
Distributed system: a sequence of local states for each process

What about in *game situations*?
Answer: a *description* of the first-order and higher-order information of the players

R. Fagin, J. Halpern and M. Vardi. *Model theoretic analysis of knowledge*. Journal of the ACM 91 (1991).

Is an Epistemic Model "Common Knowledge"?

"The implicit assumption that the information partitions...are themselves common knowledge...constitutes no loss of generality... the assertion that each individual 'knows' the knowledge operators of all individual has no real substance; it is part of the framework."

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

Is an Epistemic Model "Common Knowledge"?

"The implicit assumption that the information partitions...are themselves common knowledge...constitutes no loss of generality... the assertion that each individual 'knows' the knowledge operators of all individual has no real substance; it is part of the framework."

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

"it is an informal but *meaningful* meta-assumption....It is not trivial at all to assume it is "common knowledge" which partition every player has."

A. Heifetz. *How canonical is the canonical model? A comment on Aumann's interactive epistemology*. International Journal of Game Theory (1999).

A General Question

How many levels/states of knowledge (beliefs) are there?

A General Question

How many levels/states of knowledge (beliefs) are there?

It depends on how you count:

- Parikh and Krasucki: Countably many *levels* of knowledge
- Parikh and EP: Uncountably many levels of belief
- Hart, Heiftetz and Samet: Uncountably many *states* of knowledge

Levels of Knowledge

Fix a set of agents $\mathcal{A} = \{1, \ldots, n\}$.

$\Sigma_K = \{K_1, \ldots, K_n\}$ and $\Sigma_C = \{C_U\}_{U \subseteq \mathcal{A}}$

Level of Knowledge: $Lev_{\mathcal{M}}(p, s) = \{x \in \Sigma^* \mid \mathcal{M}, s \models xp\}$ (where $\Sigma = \Sigma_K$ or $\Sigma = \Sigma_C$).

[If Σ is a finite set, then Σ^* is the set of finite strings over Σ]
[Recall the definition of truth in a Kripke structure]

Levels of Knowledge

Fix a set of agents $\mathcal{A} = \{1, \ldots, n\}$.

$\Sigma_K = \{K_1, \ldots, K_n\}$ and $\Sigma_C = \{C_U\}_{U \subseteq \mathcal{A}}$

Level of Knowledge: $Lev_{\mathcal{M}}(p, s) = \{x \in \Sigma^* \mid \mathcal{M}, s \models xp\}$
(where $\Sigma = \Sigma_K$ or $\Sigma = \Sigma_C$).

[If Σ is a finite set, then Σ^* is the set of finite strings over Σ]
[Recall the definition of truth in a Kripke structure]

R. Parikh and P. Krasucki. *Levels of knowledge in distributed computing*. Sadhana-Proceedings of the Indian Academy of Science 17 (1992).

R. Parikh. *Levels of knowledge, games and group action*. Research in Economics 57, pp. 267 - 281 (2003).

Levels of Knowledge

A level of knowledge is simply a set of finite strings over Σ_C. Why isn't it obvious that there are *uncountably* many levels of knowledge?

Levels of Knowledge

A level of knowledge is simply a set of finite strings over Σ_C. Why isn't it obvious that there are *uncountably* many levels of knowledge?

Consider the sets:

- $L_1 = \{K_1, K_2\}$ and $L_2 = \{K_1, K_2, K_1K_2\}$
- $L_1 = \{K_1, K_3, K_1K_2K_3\}$ and $L_2 = \{K_1, K_2, K_3, K_1K_2K_3\}$

Levels of Knowledge

A level of knowledge is simply a set of finite strings over Σ_C. Why isn't it obvious that there are *uncountably* many levels of knowledge?

Consider the sets:

- $L_1 = \{K_1, K_2\}$ and $L_2 = \{K_1, K_2, K_1K_2\}$
  (*different level of knowledge*)
- $L_1 = \{K_1, K_3, K_1K_2K_3\}$ and $L_2 = \{K_1, K_2, K_3, K_1K_2K_3\}$
  (*same level of knowledge*)

Levels of Knowledge: Preliminaries

Given any two strings $x, y \in \Sigma^*$, we say x is **embeddable** in y, written $x \leq y$, if all symbols of x occur in y in the same order but not necessarily consecutively.

Levels of Knowledge: Preliminaries

Given any two strings $x, y \in \Sigma^*$, we say x is **embeddable** in y, written $x \leq y$, if all symbols of x occur in y in the same order but not necessarily consecutively. Formally, we can define $\leq$ as follows:

1. $x \leq x$ and $\epsilon \leq x$ for all $x \in \Sigma^*$
2. $x \leq y$ if there exists $x', x'', y', y'', (y, y'' \neq \epsilon)$ such that $x = x'x''$, $y = y'y''$ and $x' \leq y'$, $x'' \leq y''$.

$\leq$ is the smallest relation satisfying (1) and (2).

Levels of Knowledge: Preliminaries

Given any two strings $x, y \in \Sigma^*$, we say x is **embeddable** in y, written $x \leq y$, if all symbols of x occur in y in the same order but not necessarily consecutively. Formally, we can define $\leq$ as follows:

1. $x \leq x$ and $\epsilon \leq x$ for all $x \in \Sigma^*$
2. $x \leq y$ if there exists $x', x'', y', y'', (y, y'' \neq \epsilon)$ such that $x = x'x''$, $y = y'y''$ and $x' \leq y'$, $x'' \leq y''$.

$\leq$ is the smallest relation satisfying (1) and (2).

Example:
$aba \leq aaba$
$aba \leq abca$
$aba \not\leq aabb$

Levels of Knowledge: Preliminaries

$(X, \preceq)$ is

Levels of Knowledge: Preliminaries

$(X, \preceq)$ is

a **partial order** if $\preceq$ is reflexive, transitive and antisymmetric.

well founded if every infinite subset of X has a $(\preceq -)$minimal element.

a **linear order** if it is a partial order and all elements of X are comparable.

Levels of Knowledge: Preliminaries

$(X, \preceq)$ is

a **partial order** if $\preceq$ is reflexive, transitive and antisymmetric.

well founded if every infinite subset of X has a $(\preceq -)$minimal element.

a **linear order** if it is a partial order and all elements of X are comparable.

a **well-partial order** (WPO) if $(X, \preceq)$ is a partial order and every linear order that extends $(X, \preceq)$ (i.e., a linear order $(X, \preceq')$ with $\preceq \subseteq \preceq'$) is well-founded.

Levels of Knowledge: Preliminaries

$(X, \preceq)$ is

a **partial order** if $\preceq$ is reflexive, transitive and antisymmetric.

well founded if every infinite subset of X has a $(\preceq -)$minimal element.

a **linear order** if it is a partial order and all elements of X are comparable.

a **well-partial order** (WPO) if $(X, \preceq)$ is a partial order and every linear order that extends $(X, \preceq)$ (i.e., a linear order $(X, \preceq')$ with $\preceq \subseteq \preceq'$) is well-founded.

A set $\{a_1, a_2, \ldots\}$ of incomparable elements is a well-founded partial order but not a WPO.

Well-Partial Orders

Fact. $(X, \preceq)$ is a WPO iff $\preceq$ is well-founded and every subset of mutually incomparable elements is finite

Well-Partial Orders

Fact. $(X, \preceq)$ is a WPO iff $\preceq$ is well-founded and every subset of mutually incomparable elements is finite

Theorem (Higman). If Σ is finite, then $(\Sigma^*, \leq)$ is a WPO

G. Higman. *Ordering by divisibility in abstract algebras*. Proc. London Math. Soc. 3 (1952).

D. de Jongh and R. Parikh. *Well-Partial Orderings and Hierarchies*. Proc. of the Koninklijke Nederlandse Akademie van Wetenschappen 80 (1977).

WPO and Downward Closed Sets

Given $(X, \preceq)$ a set $A \subseteq X$ is **downward closed** iff $x \in A$ implies for all $y \preceq x$, $y \in A$.

Theorem. (Parikh & Krasucki) If Σ is finite, then there are only countably many $\leq$-downward closed subsets of Σ^* and all of them are *regular*.

Levels of Knowledge

Theorem. Consider the alphabet $\Sigma_C = \{C_U\}_{U \subseteq \mathcal{A}}$. For all strings $x, y \in \Sigma_C^*$, if $x \preceq y$ then for all pointed models $\mathcal{M}, s$, if $\mathcal{M}, s \models yP$ then $\mathcal{M}, s \models xP$.

Levels of Knowledge

Theorem. Consider the alphabet $\Sigma_C = \{C_U\}_{U \subseteq \mathcal{A}}$. For all strings $x, y \in \Sigma_C^*$, if $x \preceq y$ then for all pointed models $\mathcal{M}, s$, if $\mathcal{M}, s \models yP$ then $\mathcal{M}, s \models xP$.

Incorrect as stated, but easily fixed: every extension of a WPO is a WPO

1. $K_1 K_1 \not\preceq K_1$
2. We should have $C_U \preceq C_V$ if $U \subseteq V$.

Levels of Knowledge

Theorem. Consider the alphabet $\Sigma_C = \{C_U\}_{U \subseteq \mathcal{A}}$. For all strings $x, y \in \Sigma_C^*$, if $x \preceq y$ then for all pointed models $\mathcal{M}, s$, if $\mathcal{M}, s \models yP$ then $\mathcal{M}, s \models xP$.

Incorrect as stated, but easily fixed: every extension of a WPO is a WPO

1. $K_1 K_1 \not\preceq K_1$
2. We should have $C_U \preceq C_V$ if $U \subseteq V$.

Corollary 1. Every level of knowledge is a downward closed set.

Levels of Knowledge

Theorem. Consider the alphabet $\Sigma_C = \{C_U\}_{U \subseteq \mathcal{A}}$. For all strings $x, y \in \Sigma_C^*$, if $x \preceq y$ then for all pointed models $\mathcal{M}, s$, if $\mathcal{M}, s \models yP$ then $\mathcal{M}, s \models xP$.

Incorrect as stated, but easily fixed: every extension of a WPO is a WPO

1. $K_1 K_1 \not\preceq K_1$
2. We should have $C_U \preceq C_V$ if $U \subseteq V$.

Corollary 1. Every level of knowledge is a downward closed set.
Corollary 2. There are only countably many levels of knowledge.

Realizing Levels of Knowledge

Theorem. (R. Parikh and EP) Suppose that L is a downward closed subset of Σ_K^*, then there is a finite Kripke model $\mathcal{M}$ and state s such that $\mathcal{M}, s \models xP$ iff $x \in L$. (i.e., $L = Lev_{\mathcal{M}}(p, s)$).

States of Knowledge

S. Hart, A. Heifetz and D. Samet. *"Knowing Whether," "Knowing That," and The Cardinality of State Spaces*. Journal of Economic Theory 70 (1996).

States of Knowledge

Let W be a set of states and fix an event $X \subseteq W$.

Consider a sequence of finite boolean algebras $\mathcal{B}_0, \mathcal{B}_1, \mathcal{B}_2, \ldots$ defined as follows:
$\mathcal{B}_0 = \{\emptyset, X, \neg X, \Omega\}$
$\mathcal{B}_n = \mathcal{B}_{n-1} \cup \{K_i E \mid E \in \mathcal{B}_{n-1}, i \in \mathcal{A}\}$

The events $\mathcal{B} = \cup_{i=1,2,\ldots}\mathcal{B}_i$ are said to be **generated by** X.

States of Knowledge

Definition. Two states w, w' are **separated** by X if there exists an event E which is generated by X such that $w \in E$ and $w' \in \neg E$.

Question: How many states can be in an information structure (W, Π_1, Π_2) such that an event X separates any two of them?

States of Knowledge

Consider a K-list $(E_1, E_2, E_3, \ldots)$ of events generated by X.

We can of course, write down infinitely many infinite K-lists (uncountably many!).

States of Knowledge

Consider a K-list $(E_1, E_2, E_3, \ldots)$ of events generated by X.

We can of course, write down infinitely many infinite K-lists (uncountably many!).

Again, are they all consistent?

States of Knowledge

Consider a K-list $(E_1, E_2, E_3, \ldots)$ of events generated by X.

We can of course, write down infinitely many infinite K-lists (uncountably many!).

Again, are they all consistent?

Consider $(X, K_1X, \neg K_2K_1X, \neg K_1\neg K_2K_1X, K_2\neg K_1\neg K_2K_1X)$

States of Knowledge

Consider a K-list $(E_1, E_2, E_3, \ldots)$ of events generated by X.

We can of course, write down infinitely many infinite K-lists (uncountably many!).

Again, are they all consistent?

$(X, K_1X, \neg K_2K_1X, \neg K_1\neg K_2K_1X, K_2\neg K_1\neg K_2K_1X)$ is inconsistent.

Knowing Whether

Let $J_i E \ := \ K_i E \vee K_i \neg E$.

Lemma. Every J-list is consistent.

Theorem. (Hart, Heifetz and Samet) There exists an information structure (W, Π_1, Π_2) and an event $X \subseteq W$ such that all the states in W are separated by X and W has the cardinality of the continuum.

S. hart, A. Heifetz and D. Samet. *"Knowing Whether,", "Knowing That," and The Cardinality of State Spaces*. Journal of Economic Theory 70 (1996).

What about beliefs?

In Aumann/Kripke structures belief operators are just like knowledge operators except we replace the truth axiom/property ($K\varphi \to \varphi$) with a consistency property ($\neg B\bot$).

What about beliefs?

In Aumann/Kripke structures belief operators are just like knowledge operators except we replace the truth axiom/property ($K\varphi \rightarrow \varphi$) with a consistency property ($\neg B\bot$).

Theorem. (R. Parikh and EP) There are uncountably many levels of belief.

Returning to the Motivating Questions

- How do states of knowledge influence decisions in *game situations*?
- Can we *realize* any state of knowledge?
- What is a *state* in an epistemic model?
- Is an epistemic model *common knowledge* among the agents?

Returning to the Motivating Questions

- How do states of knowledge influence decisions in *game situations*?
- Can we *realize* any state of knowledge?
- What is a *state* in an epistemic model?
- Is an epistemic model *common knowledge* among the agents?

Returning to the Motivating Questions

- How do states of knowledge influence decisions in *game situations*?

- Can we *realize* any state of knowledge?
  It depends...

- What is a *state* in an epistemic model?
  It depends...

- Is an epistemic model *common knowledge* among the agents?
  It depends...

Ingredients of a Logical Analysis of Rational Agency

- ✓ Logics of Informational Attitudes and Informative Actions
- Logics of Motivational Attitudes (Preferences)
- Time, Action and Agency

Preference (Modal) Logics

x, y objects

$x \succeq y$: x is at least as good as y

Preference (Modal) Logics

x, y objects

$x \succeq y$: x is at least as good as y

1. $x \succeq y$ and $y \not\succeq x$ ($x \succ y$)
2. $x \not\succeq y$ and $y \succeq x$ ($y \succ x$)
3. $x \succeq y$ and $y \succeq x$ ($x \sim y$)
4. $x \not\succeq y$ and $y \not\succeq x$ ($x \perp y$)

Preference (Modal) Logics

x, y objects

$x \succeq y$: x is at least as good as y

1. $x \succeq y$ and $y \not\succeq x$ ($x \succ y$)
2. $x \not\succeq y$ and $y \succeq x$ ($y \succ x$)
3. $x \succeq y$ and $y \succeq x$ ($x \sim y$)
4. $x \not\succeq y$ and $y \not\succeq x$ ($x \perp y$)

Properties: transitivity, connectedness, etc.

Preference (Modal) Logics

Modal betterness model $\mathcal{M} = \langle W, \succeq, V \rangle$

Preference (Modal) Logics

Modal betterness model $\mathcal{M} = \langle W, \succeq, V \rangle$

Preference Modalities $\langle \succeq \rangle \varphi$: "there is a world at least as good (as the current world) satisfying φ"

$\mathcal{M}, w \models \langle \succeq \rangle \varphi$ iff there is a $v \succeq w$ such that $\mathcal{M}, v \models \varphi$

Preference (Modal) Logics

Modal betterness model $\mathcal{M} = \langle W, \succeq, V \rangle$

Preference Modalities $\langle \succeq \rangle \varphi$: "there is a world at least as good (as the current world) satisfying φ"

$\mathcal{M}, w \models \langle \succeq \rangle \varphi$ iff there is a $v \succeq w$ such that $\mathcal{M}, v \models \varphi$

$\mathcal{M}, w \models \langle \succ \rangle \varphi$ iff there is $v \succeq w$ and $w \not\succeq v$ such that $\mathcal{M}, v \models \varphi$

Preference (Modal) Logics

1. $\langle \succ \rangle \varphi \rightarrow \langle \succeq \rangle \varphi$
2. $\langle \succeq \rangle \langle \succ \rangle \varphi \rightarrow \langle \succ \rangle \varphi$
3. $\varphi \wedge \langle \succeq \rangle \psi \rightarrow (\langle \succ \rangle \psi \vee \langle \succeq \rangle (\psi \wedge \langle \succeq \rangle \varphi))$
4. $\langle \succ \rangle \langle \succeq \rangle \varphi \rightarrow \langle \succ \rangle \varphi$

Theorem The above logic (with Necessitation and Modus Ponens) is sound and complete with respect to the class of preference models.

J. van Benthem, O. Roy and P. Girard. *Everything else being equal: A modal logic approach to* ceteris paribus *preferences*. JPL, 2008.

Preference Modalities

$\varphi \geq \psi$: the state of affairs φ is at least as good as ψ
(ceteris paribus)

G. von Wright. *The logic of preference*. Edinburgh University Press (1963).

Preference Modalities

$\varphi \geq \psi$: the state of affairs φ is at least as good as ψ (ceteris paribus)

G. von Wright. *The logic of preference*. Edinburgh University Press (1963).

$\langle \Gamma \rangle^{\leq} \varphi$: φ is true in "better" world, *all things being equal*.

J. van Benthem, O. Roy and P. Girard. *Everything else being equal: A modal logic approach to* ceteris paribus *preferences*. JPL, 2008.

All Things Being Equal...

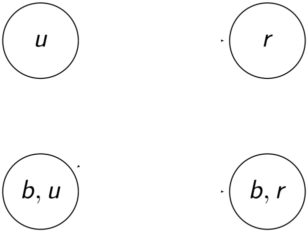

All Things Being Equal...

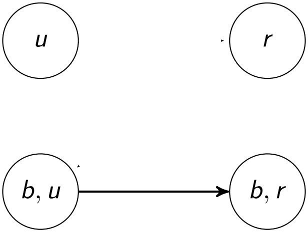

- With boots (b), I prefer my raincoat (r) over my umbrella (u)

All Things Being Equal...

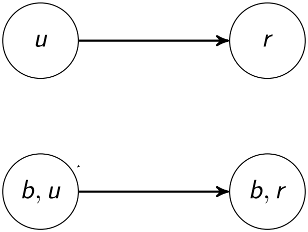

- With boots (b), I prefer my raincoat (r) over my umbrella (u)
- Without boots ($\neg b$), I also prefer my raincoat (r) over my umbrella (u)

All Things Being Equal...

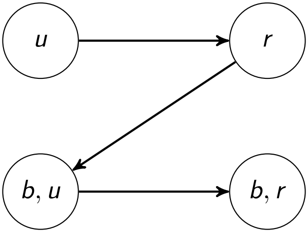

- With boots (b), I prefer my raincoat (r) over my umbrella (u)
- Without boots ($\neg b$), I also prefer my raincoat (r) over my umbrella (u)
- But I do prefer an umbrella and boots over a raincoat and no boots

All Things Being Equal...

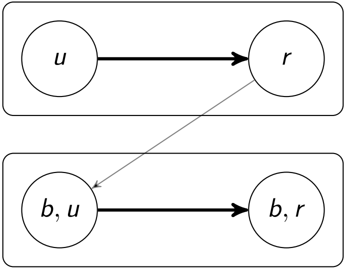

All things being equal, I prefer my raincoat over my umbrella

All Things Being Equal...

Let Γ be a set of (preference) formulas. Write $w \equiv_\Gamma v$ if for all $\varphi \in \Gamma$, $w \models \varphi$ iff $v \models \varphi$.

1. $\mathcal{M}, w \models \langle\Gamma\rangle\varphi$ iff there is a $v \in W$ such that $w \equiv_\Gamma v$ and $\mathcal{M}, v \models \varphi$.
2. $\mathcal{M}, w \models \langle\Gamma\rangle^{\leq}\varphi$ iff there is a $v \in W$ such that $w(\equiv_\Gamma \cap \leq)v$ and $\mathcal{M}, v \models \varphi$.
3. $\mathcal{M}, w \models \langle\Gamma\rangle^{<}\varphi$ iff there is a $v \in W$ such that $w(\equiv_\Gamma \cap <)v$ and $\mathcal{M}, v \models \varphi$.

All Things Being Equal...

Let Γ be a set of (preference) formulas. Write $w \equiv_\Gamma v$ if for all $\varphi \in \Gamma$, $w \models \varphi$ iff $v \models \varphi$.

1. $\mathcal{M}, w \models \langle\Gamma\rangle\varphi$ iff there is a $v \in W$ such that $w \equiv_\Gamma v$ and $\mathcal{M}, v \models \varphi$.
2. $\mathcal{M}, w \models \langle\Gamma\rangle^{\leq}\varphi$ iff there is a $v \in W$ such that $w(\equiv_\Gamma \cap \leq)v$ and $\mathcal{M}, v \models \varphi$.
3. $\mathcal{M}, w \models \langle\Gamma\rangle^{<}\varphi$ iff there is a $v \in W$ such that $w(\equiv_\Gamma \cap <)v$ and $\mathcal{M}, v \models \varphi$.

Key Principles:

- $\langle\Gamma'\rangle\varphi \rightarrow \langle\Gamma\rangle\varphi$ if $\Gamma \subseteq \Gamma'$
- $\pm\varphi \wedge \langle\Gamma\rangle(\alpha \wedge \pm\varphi) \rightarrow \langle\Gamma \cup \{\varphi\}\rangle\alpha$

All Things Being Equal...

Let Γ be a set of (preference) formulas. Write $w \equiv_\Gamma v$ if for all $\varphi \in \Gamma$, $w \models \varphi$ iff $v \models \varphi$.

1. $\mathcal{M}, w \models \langle\Gamma\rangle\varphi$ iff there is a $v \in W$ such that $w \equiv_\Gamma v$ and $\mathcal{M}, v \models \varphi$.
2. $\mathcal{M}, w \models \langle\Gamma\rangle^{\leq}\varphi$ iff there is a $v \in W$ such that $w(\equiv_\Gamma \cap \leq)v$ and $\mathcal{M}, v \models \varphi$.
3. $\mathcal{M}, w \models \langle\Gamma\rangle^{<}\varphi$ iff there is a $v \in W$ such that $w(\equiv_\Gamma \cap <)v$ and $\mathcal{M}, v \models \varphi$.

Key Principles:

- $\langle\Gamma'\rangle^{\leq}\varphi \rightarrow \langle\Gamma\rangle^{\leq}\varphi$ if $\Gamma \subseteq \Gamma'$
- $\pm\varphi \wedge \langle\Gamma\rangle^{\leq}(\alpha \wedge \pm\varphi) \rightarrow \langle\Gamma \cup \{\varphi\}\rangle^{\leq}\alpha$

To be continued....