Syntactic Epistemic Logic

Sergei Artemov

Graduate Center CUNY

September 27, 2014

Abstract

We offer a paradigm shift in Epistemic Logic: **to view an epistemic scenario as specified syntactically by a set of formulas** in an appropriate extension of epistemic modal logic.

We make a case that the alternative approach to specify an epistemic scenario as a Kripke/Aumann model is unnecessarily restrictive. Some/many scenarios that admit natural syntactic formalization do not have independent model characterizations.

On the other hand, the syntactic approach is inclusive, e.g., a semantic specification by a finite Kripke/Aumann model yields a linear-time decidable syntactic description.

Formal syntactic specifications can be studied with the entire spectrum of tools, including deduction and semantic modeling.

Disclaimer

The observations upon which we base our proposal are mostly commonplace for a professional logician. However, what we want to promote is changing the way logicians and experts in epistemic-related applications, first of all in Game Theory, **specify/formalize** epistemic scenarios:

1. don't view Kripke/Aumann specifications as universal - syntactic specs are more general;
2. if a scenario is originally described syntactically and formalized as a model, think of justification, e.g., prove completeness.

By no means we want to discriminate against the semantic approach which defines epistemic scenarios as Kripke/Aumann structures; constructive semantic specifications are accepted.

About the title

The name *Syntactic Epistemic Logic* was suggested by Robert Aumann who pointed to the conceptual and technical gap between the syntactic character of game descriptions and the predominantly semantical way of analyzing games via possible world/partition models.

Hidden dangers of the semantic approach

We adopt the aforementioned view that the initial description $\mathcal{I}$ of an epistemic situation is syntactic and informal in a natural language. The long-standing tradition in epistemic logic and game theory is "given $\mathcal{I}$, proceed to a specific epistemic model $\mathcal{M}_{\mathcal{I}}$ and make the latter a mathematical definition of $\mathcal{I}$":

$$\textit{informal syntactic description } \mathcal{I} \;\; \Rightarrow \;\; \textit{`natural' model } \mathcal{M}_{\mathcal{I}}. \qquad (1)$$

There are hidden dangers in this process: a syntactic description $\mathcal{I}$ may have multiple models and picking one of them (especially declaring it common knowledge) requires justification. Furthermore, if we seek an exact specification, then some/many scenarios that have natural syntactic formalization do not have epistemically acceptable model descriptions at all.

Going syntactic: In the beginning was the Word

Through the framework of *Syntactic Epistemic Logic*, SEL, we suggest making **the syntactic logic formalization $\mathcal{S}_{\mathcal{I}}$ a formal definition of the situation described by $\mathcal{I}$**:

$$\textit{description } \mathcal{I} \;\Rightarrow\; \textit{formalization } \mathcal{S}_{\mathcal{I}} \Rightarrow \textit{all its models } \mathcal{M}_{\mathcal{S}}. \quad (2)$$

The first step from $\mathcal{I}$ to $\mathcal{S}_{\mathcal{I}}$ is normally straightforward and deterministic, barring ambiguities of $\mathcal{I}$.

Step 2 from $\mathcal{S}_{\mathcal{I}}$ to $\mathcal{M}_{\mathcal{S}}$'s is mathematically rigorous, since $\mathcal{S}_{\mathcal{I}}$ has a well-defined class of models.

Approach (2) is scientific and, as we argue, encompasses a broader class of epistemic scenarios than the semantic approach (1).

Basics of epistemic logic and its models

The logic language is augmented by modalities $\mathbf{K}_1, \mathbf{K}_2, \ldots$, for agents' knowledge. Models are sets of possible worlds with indistinguishability relations $R_1, R_2, \ldots$, and truth values of atoms at each world. 'F holds at u' ($u \Vdash F$) respects Booleans and

$$u \Vdash \mathbf{K}_i F \quad \textit{iff} \quad v \Vdash F \textit{ for each state } v \textit{ s.t. } uR_i v.$$

Example: states $\{u, v, w\}$, R_1 - the solid arrow, R_2 - dotted.

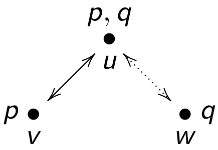

$u \Vdash \mathbf{K}_1 p$ and $u \Vdash \mathbf{K}_2 q$, but not vice versa: $u \nVdash \mathbf{K}_1 q$ and $u \nVdash \mathbf{K}_2 p$.

Basics of epistemic logic and its models

Let Γ be a set of epistemic formulas. A **model** of Γ is an epistemic structure $\mathcal{M}$ and a state ω s.t. all formulas from Γ are true at ω:

$$\mathcal{M}, \omega \Vdash \Gamma.$$

A formula F **follows semantically** from Γ, $\Gamma \models F$, if F holds in each model of Γ. A well-known fact: **Completeness Theorem**

$$\Gamma \vdash F \quad \Leftrightarrow \quad \Gamma \models F.$$

This has been used by some to claim the equivalence of the syntactic and semantic approaches in epistemology, in particular to justify specifying epistemic scenarios semantically by an epistemic model structure. We will challenge these claims and show the limitations of semantic specifications.

Epistemic states and canonical models

Completeness Theorem claims that if $\Gamma \not\vdash F$ then there is a model $\mathcal{M}, \omega$ in which F is false. Where does this model come from?

In any model $\mathcal{M}, \omega$, the set of truths $\mathcal{T}$ contains Γ and is **maximal**, i.e., for each formula F, $\mathcal{T}$ contains F or contains $\neg F$. This observation suggests the notion of

epistemic state = maximal consistent extension of Γ.

A comprehensive "canonical" model of Γ consists of all possible epistemic states over Γ and typically has continuum elements. Epistemic relations are also defined on the basis of what is known at each state: for maximal consistent α and β, $\alpha R_i \beta$ iff for each F

$$\mathbf{K}_i F \in \alpha \quad \Rightarrow \quad F \in \beta.$$

The good and bad about canonical models

The good: the completeness claim is immediate: if Γ does not prove F, then $\Gamma + \{\neg F\}$ is consistent and hence can be extended to a maximal consistent set (epistemic state) in which F is false.

The bad: a canonical model is not an independently defined semantic structure for specifying knowledge assertions. On the contrary, states and relations of the canonical model are reverse engineered from syntactic data of what is known at each world.

Conceptually, the canonical model $\mathcal{M}(\Gamma)$ of an epistemic scenario Γ **cannot be used as a semantic definition of** Γ just because Γ itself is needed to define $\mathcal{M}(\Gamma)$. *We do not predict the weather for yesterday using yesterday's meteorological readings.*

Canonical models are typically too big to be known

In some epistemic contexts, e.g., in Aumann's partition models, there is a *common knowledge of the model* requirement (which is justified if the model is the everyone's source of epistemic data).

We argue that some/many epistemic scenarios with reasonable syntactic descriptions Γ have canonical models $\mathcal{M}(\Gamma)$ with continuum epistemic states. Such models cannot be known and such scenarios have no satisfactory semantic characterizations.

In summary, the canonical model $\mathcal{M}(\Gamma)$ is a derivative of Γ. Furthermore, if $\mathcal{M}(\Gamma)$ is generic (continuum states, unknowable) there are no reasons to consider $\mathcal{M}(\Gamma)$ as a primary semantic characterization of the epistemic problem.

For some scenarios, the semantic characterization works

The situation is quite different if the canonical model $\mathcal{M}(\Gamma)$ is proved to collapse into a reasonable finite model $\mathcal{M}'(\Gamma)$ (it will be the case with the paradigmatic Muddy Children problem): then $\mathcal{M}'(\Gamma)$ is a helpful semantic characterization of Γ, e.g., provability in Γ is linear-time decidable.

Baby examples first

Example: two agents and two propositional variables p_1 and p_2.

1. $\Gamma = \{p_1 \wedge p_2\}$, i.e., both atoms are true. The corresponding canonical model has continuum-many states: there are infinitely many sufficiently independent higher-order epistemic assertions.

2. $\Gamma = \{\mathbf{C}(p_1 \wedge p_2)\}$, i.e., it is common knowledge that both atoms are true. There is only one epistemic state at which both p_1 and p_2 hold (hence are common knowledge).

Muddy Children: informal description

Consider the Muddy Children puzzle, which is formulated syntactically and can be formalized in multi-agent epistemic logic.

> *A group of n children meet their father after playing in the mud. Their father notices that $k > 0$ of the children have mud on their foreheads. Each child sees everybody else's foreheads, but not his own. The father says: "some of you are muddy," then says: "Do any of you know that you have mud on your forehead? If you do, raise your hand now." No one raises his hand. The father repeats the question, and again no one moves. After exactly k repetitions, all children with muddy foreheads raise their hands simultaneously.*

Why?

Muddy Children: syntactic formalization

This situation can be described in epistemic logic with atomic propositions $m_1, m_2, \ldots, m_n$ with m_i stating that child i is muddy, and modalities $\mathbf{K}_1, \mathbf{K}_2, \ldots, \mathbf{K}_n$ for the children' knowledge.

In addition to general epistemic logic principles, the scenario description includes the following set MC_n of assumptions:

1. Knowing about the others:

$$\bigwedge_{i \neq j} [\mathbf{K}_i(m_j) \vee \mathbf{K}_i(\neg m_j)].$$

2. Not knowing about himself:

$$\bigwedge_{i=1,\ldots,n} [\neg \mathbf{K}_i(m_i) \wedge \neg \mathbf{K}_i(\neg m_i)].$$

Muddy Children: syntactic solution

Consider the case $n = k = 2$, i.e., two children, both muddy. Here is an informal solution of the problem.

> *After father's announcement "some of you are muddy," if a child sees another child not muddy, he knows that he himself is muddy. This argument is known to everybody, and since both children announce that they did not know that they were muddy, both figure out that they are muddy and raise their hands in the second round.*

This reasoning is quite rigorous and can be itself directly formalized within an appropriate modal epistemic logic.

Muddy Children: semantic solution

In a model-theoretical solution, the set of assumptions MC_n is replaced by an *ad hoc* Kripke model: n-dimensional cube Q_n. Again, consider $n = 2$.

Logical possibilities for the truth value combinations of (m_1, m_2): (0,0), (0,1), (1,0), and (1,1) are declared epistemic states. There are two indistinguishability relations denoted by the dotted arrows (for agent 1) and the solid arrows (for agent 2).

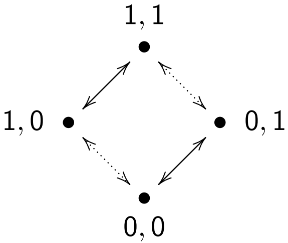

Model Q_2

Muddy Children: semantic solution

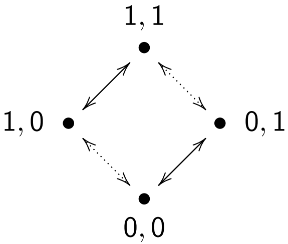

Model Q_2

It is easy to check that conditions 1 (knowing about the other) and 2 (not knowing about himself) hold at each node of this model. Furthermore, Q_2 is assumed to be commonly known to the agents.

Muddy Children: semantic solution

After the father publicly announces $m_1 \vee m_2$, node $0, 0$ is no longer possible:

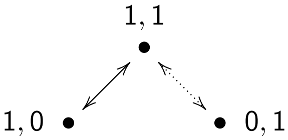

Model $\mathcal{M}_1$.

$\mathcal{M}_1$ now becomes common knowledge. Both children realize that in 1,0 child 2 would know whether he is muddy (no other 2-indistinguishable worlds), and in 0,1, child 1 would know.

Muddy Children: semantic solution

After both children answer “No” to the question of whether they know what is on their foreheads, worlds 1,0 and 0,1 are no longer possible, and each child eliminates them from the set of possible worlds. The only remaining logical possibility here is

$$1, 1$$

•

Model $\mathcal{M}_2$.

Now both children know that their foreheads are muddy.

What is missing in the semantic solution?

This semantic solution starts with adopting a model, Q_n, as an equivalent of a theory, MC_n. Is such a reduction justified?

One needs to show that MC_n and Q_n describe the same set of truths (so far we can claim only that all truths of MC_n hold in Q_n). In the case of MC_n and Q_n this can be done.

Let u be a vertex in Q_n. We define its formal representation

$$\pi(u) = \bigwedge\{m_i \mid u \Vdash m_i\} \wedge \bigwedge\{\neg m_j \mid u \nVdash m_j\}.$$

Proposition 1. (Completeness of MC_n):

F holds at u in Q_n $\quad\Leftrightarrow\quad$ MC_n *proves* $\pi(u) \rightarrow F$.

Corollary. *There are 2^n epistemic states in MC_n; they correspond to nodes of Q_n which is therefore a canonical model for MC_n.*

Muddy Children scenario is a lucky exception

Accidentally, in this case, MC_n, picking one "natural model" (here Q_n) can be justified: Proposition 1 (easy, but not entirely trivial) states that MC_n is complete w.r.t. Q_n, hence each logical property of Q_n is derivable in MC_n.

However, in a general setting, the approach

given a syntactic description pick a "natural model"

is intrinsically flawed: a completeness analysis is required.

Why is assuming Q_n for Muddy Children a big deal?

A possible (and observed) reaction from an epistemic logician on the criticism that Q_n was adopted as *The Model* of MC_n without a completeness analysis of MC_n is

> *It is not much to assume that an agent can figure out that the logical possibilities correspond to the vertices of Q_n, e.g., for n=2, they are (0,0), (0,1), (1,0), and (1,1).*

This argument only goes halfway: **it does not explain why a combination q of truth values of atoms determine, in MC_n, truth values of any relevant epistemic sentence**. Without this, we cannot claim that q is an epistemic state.

Let us consider an example when it is not.

Muddy Children 'lite'

Consider a simplified Muddy Children scenario, MClite_2, in which condition (2) "not knowing about himself" is omitted:

> *Two children have muddy foreheads and each child sees the other child's forehead. The father announces publicly "some of you are muddy."*
>
> *The father then says: "Do any of you know that you have mud on your forehead? If you do, raise your hand now." No one raises his hand. The father repeats the question, and both children raise their hands simultaneously.*

What is the natural epistemic model of the starting configuration?

Muddy Children 'lite,' epistemic models

Good old Q_2 is certainly a model (fewer conditions to check than for MC_2).

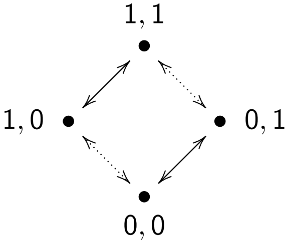

Model Q_2

Muddy Children 'lite,' epistemic models

Here is another model, not equivalent to Q_2:

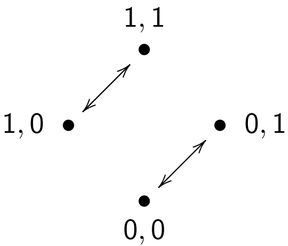

Model $\mathcal{M}_3$.

Muddy Children 'lite,' epistemic models

And another:

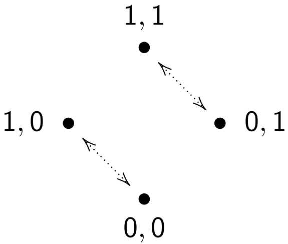

Model $\mathcal{M}_4$.

Muddy Children 'lite,' epistemic models

And even this:

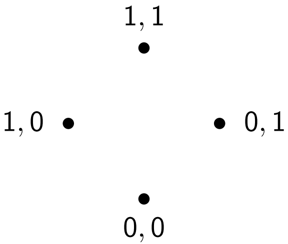

Model $\mathcal{M}_5$.

None of these models alone adequately represent MClite_2.

It is not just about nondeterministic edges

Incidentally, multiple models here appear not just because some edges in Q_2 are not specified in MClite_2. The problem is deeper: even within one model, **truth values of atomic propositions do not necessarily determine an epistemic state**.

Example. MClite_2 holds at each world of $\mathcal{M}_6$

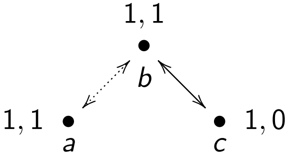

Model $\mathcal{M}_6$

which has different worlds a and b with the same propositional component 1,1: $a \Vdash \mathbf{K}_2 m_2$ and $b \Vdash \neg\mathbf{K}_2 m_2$.

MClite$_2$ has continuum epistemic states

Proposition 2. *MClite$_2$ has continuum different epistemic states.*

Proof idea. Even though m_1 holds, $\mathbf{K}_1(m_1)$ is independent. Furthermore, the second-order 2-knowledge $\mathbf{K}_2\mathbf{K}_1(m_1)$ and $\mathbf{K}_2\neg\mathbf{K}_1(m_1)$ is also independent, as are similar higher-order alternating epistemic assertions concerning m_1. Hence continuum many choices, all consistent and pairwise incompatible. The rigorous proof involves some model reasoning.

As we see, the only available semantic description of MClite$_2$ is a generic canonical model which is a non-constructive derivative of the syntactic description Γ and not very helpful.

Muddy Children 'lite,' syntactic solution

The natural logic solution, nevertheless, stands:

> *After father's announcement "some of you are muddy," if a child sees another child not muddy, he knows that he himself is muddy. This argument is known to everybody, and since both children announce that they did not know that they were muddy, both figure out that they are muddy and raise their hands in the second round.*

This reasoning does not use condition (2) "not knowing about himself" and hence is good for MClite_2 as well.

What is the true meaning of the model solution?

1. The model solution w/o completeness analysis uses a strong additional assumption (common knowledge) of a specific model and hence does not resolve the original Muddy Children puzzle; it corresponds to a different scenario, e.g.,

> *A group of robots programmed to reason about model Q_n meet their programmer after playing in the mud. ...*

2. One could argue that the model solution actually codifies a deductive solution in the same way that geometric reasoning is merely a visualization of a rigorous derivation in some sort of axiom system for geometry. This is a valid point which can be made scientific within the framework of Syntactic Epistemic Logic.

Aumann's extensive games

Aumann's definition of an extensive game is based on the notion of a partition structure, essentially equivalent to S5 Kripke models with an extra condition that the model should itself be common knowledge. This postulates a convenient framework for reasoning about games.

We argue, however, that this definition of a game is too restrictive and needs to be extended.

Partition models

Aumann defines an extensive game as a game tree and a partition structure $\mathcal{A}$ which are commonly known. It is presumed that $\mathcal{A}$ codifies possible epistemic states of the game, i.e., what is known and what is not known to the players.

Available strategies are represented by atomic propositions and epistemic conditions are formulas in the multi-agent epistemic logic over these propositions. The principal epistemic reading of $\mathbf{K}_i F$ is:

$\mathbf{K}_i F$ *holds at* ω *iff* F *holds at any* ω' *undistinguishable from* ω.

Common knowledge of rationality yields an exact model

Consider Game 1 with the game tree shown and with the standard assumption of *common knowledge of the game and rationality*

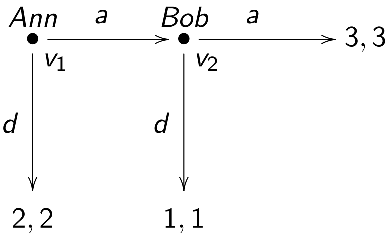

of players, CKGR. Since Bob is rational, he plays across. Ann knows that Bob is rational; she anticipates Bob's move and herself moves across, hence the Backward Induction solution (a, a) which is common knowledge.

Assume the logic language of this game consists of propositional variables a_A and a_B symbolizing Ann's and Bob's moves across, logical connectives, and knowledge modalities $\mathbf{K}_A$ and $\mathbf{K}_B$. The syntactic formalization can be reduced to the set of formulas

$$S = \{\mathbf{C}a_B, \mathbf{C}(\mathbf{K}_A a_B \rightarrow a_A)\},$$

where $\mathbf{C}$ is the common knowledge modality.

The corresponding Kripke model $(W, R_A, R_B, \Vdash)$ has one node (a, a), R_A and R_B are reflexive, a_A and a_B hold at (a, a).

Model $\mathcal{M}_7$

A game that does not have an exact partition model

Consider Game 2 that features the same players Ann and Bob, and the game tree as Game 1. The difference is in the epistemic conditions: *Ann is rational and knows that Bob is rational*. Note that the game does not specify whether Bob knows that Ann is rational, not to mention higher-order epistemic assertions of type "Ann knows that Bob does not know that she is rational," etc.

The solution of Game 2 is the same: Ann and Bob play across. This can be naturally established by a straightforward syntactic formalization of Game 2, followed by an easy logical reasoning.

Syntactic formalization and solution of Game 2

1. *Bob is rational* is formalized as a_B (moves for a higher payoff); *Ann knows that Bob is rational* is formalized as

$$\mathbf{K}_A(a_B);$$

2. *Ann is rational* is formalized as *if Ann knows that Bob plays across, she plays across*:

$$\mathbf{K}_A(a_B) \to a_A;$$

3. *knowing their moves*:

$$a_i \to \mathbf{K}_i(a_i), \quad \neg a_i \to \mathbf{K}_i(\neg a_i), \quad i \in \{A, B\}.$$

The formalization of Game 2 is then the set $\Gamma_2 = \{1, 2, 3\}$. Obviously, Γ_2 proves a_A, a_B, $\mathbf{K}_A(a_B)$, but $\mathbf{K}_B(a_A)$ is independent.

An attempt to find a model formulation for Game 2

A natural attempt to formalize Game 2 by a model could be $\mathcal{M}_8$. The dotted arrow is indistinguishability for Bob, both nodes are R_A- and R_B-reflexive.

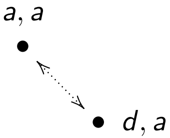

Model $\mathcal{M}_8$

At (a, a), Ann knows that Bob is playing across, so $\mathcal{M}_8$ at (a, a) is a model of Game 2. However, $\neg\mathbf{K}_B(a_A)$ holds at (a, a) but does not follow from Game 2. So $\mathcal{M}_8$ is not an exact model of Game 2.

Model $\mathcal{M}_8$ can be represented as an Aumann structure $\mathcal{A}_8$ that has two epistemic states,

$$\Omega = \{\omega_1, \omega_2\}$$

where ω_1 corresponds to profile (a, a) and ω_2 - to profile (d, a). Partitions of Ω for Ann and Bob are

$$\mathcal{K}_A = \{\{\omega_1\}, \{\omega_2\}\}, \qquad \mathcal{K}_B = \{\{\omega_1, \omega_2\}\},$$

and the real state is ω_1.

In $\mathcal{K}_A$, both partition cells are singletons, hence Ann knows Bob's move. In $\mathcal{K}_B$, states ω_1 and ω_2 are in the same cell and hence are indistinguishable for Bob who does not know Ann's move.

Therefore, $\neg\mathbf{K}_B(a_A)$ holds in $\mathcal{A}_8$, but does not follow from Game 2.

There is no ‘knowable’ semantic formulation of Game 2

Proposition 3. *Game 2 (formalized by Γ_2) has continuum-many different epistemic states*.

Proof idea. The proof plays with higher-order epistemic assertions that are independent from Game 2. Some model combinatorics is needed to make this observation precise.

Corollary. *Game 2 cannot be represented by a commonly known partition model*.

Syntactic formalization is there anyway

A step from an informal syntactic description of an epistemic scenario to its **formalized syntactic version** Γ does not appear to have built-in foundational logical problems of connecting syntactic and semantics realms.

If a semantic characterization of Γ is given by a generic canonical model $\mathcal{M}(\Gamma)$, then Γ remains the original characterization and $\mathcal{M}(\Gamma)$ its non-constructive unknowable derivative.

If the scenario can be naturally formalized by a manageable model, this is a win-win situation. A finite model automatically defines a syntactic set of true formulas Γ which is linear-time decidable (an impressive speed-up compared to PSPACE-completeness of the multi-agent epistemic logic itself).

“Dark matter” of the epistemic universe

How typical are manageable semantic specifications? We argue that this is an exception rather than the rule.

> *It appears that unless common knowledge of the basic assumptions is postulated in Γ, independent higher-order epistemic assertions supply enough building material for continuum many different epistemic states and render semantic specifications nonconstructive/unknowable.*

Epistemic conditions more flexible than CKGR (mutual knowledge of rationality, asymmetric epistemic assumptions, as in Game 2, etc.) lead to generic/unknowable canonical models.

These cases are like the“dark matter” of epistemic universe: they are everywhere, but cannot be visualized by knowable partition models. **The semantic approach does not recognize this “dark matter”: SEL deals with it syntactically**.

Why does the semantic approach often work?

An interesting question is why the semantic approach, despite its aforementioned shortcomings, produces correct answers in many situations. We see several reasons for this.

1. The (once) standard *common knowledge of the game and rationality* assumption yields common knowledge of the Backward Induction solution. For such games, their epistemic structure reduces to non-problematic singleton models, no "dark matter."

> *However, game theorists consider CKGR too restrictive: players might not have complete and equal information about the game and each other, there might be a certain amount of ignorance and/or secrecy, etc. In these cases, CKGR does not hold.*

Why does the semantic approach often work?

2. Pragmatic self-limitation. Given an informal description of a game G, we intuitively seek a solution that logically follows from G. Even if we skip the formalization of G and pick its 'natural' model $\mathcal{M}(G)$, not necessarily capturing the whole G, we try not to use features of the model that are not supported by G. If we conclude a property P by such restricted reasoning about the model, then P indeed logically follows from G.

Such an *ad hoc* pragmatic approach can be made scientific within the framework of Syntactic Epistemic Logic.

> *This resembles Geometry, in which we reason about triangles, circles, etc., but on the background have a rigorous system of postulates and we are trained not to go beyond there postulates.*

What do we gain by going syntactic?

The Syntactic Epistemic Logic suggestion: make the syntactic logic formalization of an epistemic scenario its formal specification. What do we gain by going syntactic?

SEL does not lose any of the advantages of the semantic methods and offers a cure for two principal weaknesses of the latter:

1. SEL provides a scientific framework for resolving the tension between a syntactic description and its hand-picked model: formalize the former and prove completeness.

2. SEL suggests a way to handle "dark matter" scenarios with non-constructive/unknowable models which, however, have reasonable axiomatic descriptions (MClite, Game 2, etc.).

Extended definition of an extensive game

Within Syntactic Epistemic Logic, we offer a new definition of an extensive game which is more general than Aumann's partition model definition:

> *a game tree supplied by a syntactic description* Γ *of epistemic conditions in an appropriate extension of multi-agent epistemic logic.*

Again, numerous games with "dark matter" epistemic conditions (generic non-constructive and unknowable canonical models) can now be formalized (cf. Game 2) and studied.

Some practical implications

If an epistemic scenario is given syntactically, but formalized by an epistemic model, it makes sense to examine its syntactic formalization as well and try to establish their equivalence.

A broad class of epistemic scenarios underdefine higher epistemic assertions (individual knowledge, mutual and limited-depth knowledge, partial knowledge, etc.), have continuum epistemic states and no satisfactory partition models. However, if such a scenario allows an adequate syntactic formulation, it can be handled with the entire spectrum of mathematical tools.

Since the basic object in SEL is a syntactic description Γ of an epistemic scenario rather than a specific model, there is room for a new syntactic theory of updates and belief revision.