

Towards Syntactic Epistemic Logic

Sergei Artemov*

The Graduate Center, The City University of New York

365 Fifth Avenue, New York City, NY 10016, USA

sartemov@gc.cuny.edu

Abstract. Traditionally, Epistemic Logic represents epistemic scenarios using a single model. This, however, covers only complete descriptions that specify truth values of *all* assertions. Indeed, many—and perhaps most—epistemic descriptions are not complete. Syntactic Epistemic Logic, SEL, suggests viewing an epistemic situation as a set of syntactic conditions rather than as a model. This allows us to naturally capture incomplete descriptions; we discuss a case study in which our proposal is successful. In Epistemic Game Theory, this closes the conceptual and technical gap, identified by R. Aumann, between the syntactic character of game-descriptions and semantic representations of games.

1. Introduction

In this paper, we argue for a paradigm shift in the way that logic and epistemic-related applications – in particular, game theory – specify epistemic scenarios.¹ Given a verbal description of a situation, a typical epistemic user cherry-picks a “natural model” (Kripke or Aumann’s) and then regards it as a formalization of the original description. This approach carries with it two fundamental deficiencies:

- I. It covers only complete descriptions, whereas many (intuitively most) epistemic situations are partially described and cannot be adequately specified by a single model.²

*Address for correspondence: The Graduate Center, The City University of New York, 365 Fifth Avenue, New York City, NY 10016, USA.

¹The preliminary version of this paper was delivered as an invited talk at the 15th LMPS Congress in 2015 [3].

²Epistemic logicians have been mostly aware of (I) but this did not stop the wide spread culture of identifying an epistemic scenario with a single Kripke model (or Aumann structure in Game Theory).

II. The traditional epistemic reading of Kripke/Aumann models requires common knowledge of the model which restricts their generality and utility even further.

1.1. Overspecification

A typical case of (I) is the overspecification problem. Consider the following description:

A tossed coin lands heads up. Alice sees the coin, Bob does not. (1)

Students in an epistemic logic class normally produce a Kripke S5-model of this situation as in Figure 1.




Figure 1. Model \mathcal{M}_1 .

In this model, there are two possible worlds 1 and 2, arrows represent indistinguishability relations R_A and R_B between worlds, h is a propositional letter for “heads,” and node 1 represents the real world at which h holds.

\mathcal{M}_1 is a model of (1) which, however, overspecifies (1): in this model there are propositions which are true but do not follow from (1), e.g.,

- $\mathbf{K}_A \neg \mathbf{K}_B h$ - Alice knows that Bob does not know h ;³
- $\mathbf{K}_B (\mathbf{K}_A h \vee \mathbf{K}_A \neg h)$ - Bob knows that Alice knows whether h ;
- *etc.*

We will see in Section 4 that scenario (1) “as is” does not have a single-model specification at all.

In a situation in which an epistemic scenario is described syntactically but formalized as a model, a completeness analysis relating these two modes is required. For example, the Muddy Children puzzle is given syntactically but then presented as a model tacitly presumed to be commonly known (cf. [14, 16, 17, 18, 19]). In Section 5, we show that this choice of a specifying model can be justified. However, the Muddy Children case is a fortuitous exception: see Sections 5 and 6 for more epistemic scenarios without single model specifications.

Existing approaches to mitigate overspecification include

- Supervaluations: given a syntactically defined situation \mathcal{S} , assume

“F holds in \mathcal{S} ” iff “F is true in all models of \mathcal{S} .”

This approach has been dominant in mathematical logic with formal theories as “situations,” and it manifests itself in Gödel’s Completeness Theorem.

³ \mathbf{K}_A and \mathbf{K}_B are knowledge modalities for Alice and Bob.

- Non-standard truth values: Kleene three-valued logic or other, more exotic ways of defining truth values. This approach has generated mathematically attractive models, but it has neither dethroned the supervaluation tradition in mathematical logic, nor changed the ill-founded “natural model culture” in epistemology.

Here we explore the supervaluation approach in epistemology by representing epistemic scenarios in a logical language syntactically and considering the whole class of the corresponding models, not just one cherry-picked model. This also eliminates problem (II).

2. What is Syntactic Epistemic Logic?

The name *Syntactic Epistemic Logic* was suggested by Robert Aumann (cf. [9]) who identified the conceptual and technical gap between the syntactic character of game descriptions and the predominantly semantic way of analyzing games via relational/partition models.

Suppose the initial description \mathcal{I} of an epistemic situation is syntactic in a natural language. The long-standing tradition in epistemic logic and game theory is then to proceed to a specific epistemic model $\mathcal{M}_{\mathcal{I}}$, and take the latter as a mathematical definition of \mathcal{I} :

$$\text{informal description } \mathcal{I} \Rightarrow \text{“natural model” } \mathcal{M}_{\mathcal{I}}. \quad (2)$$

Hidden dangers lurk within this process: a syntactic description \mathcal{I} may have multiple models and picking one of them (especially declaring it common knowledge) is not generally sound. Furthermore, if we seek an exact specification, then only deductively complete scenarios can be represented (cf. Theorem 4.3). Epistemic scenarios outside this group, which include situations with asymmetric and less-than-common knowledge (e.g., mutual knowledge) of conditions, do not have single-model presentations, but can be specified and handled syntactically.

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

$$\text{description } \mathcal{I} \Rightarrow \text{syntactic formalization } \mathcal{S}_{\mathcal{I}} \Rightarrow \text{all of } \mathcal{S}_{\mathcal{I}}\text{'s models}. \quad (3)$$

The first step from \mathcal{I} to $\mathcal{S}_{\mathcal{I}}$ is formalization and it has its own subtleties which we will not analyze here.

The SEL approach (3), we argue, encompasses a broader class of epistemic scenarios than a semantic approach (2). In this paper, we provide motivations and sketch basic ideas of Syntactic Epistemic Logic. Specific suggestions of general purpose formal systems is a work in progress, cf. [4].

SEL provides a more balanced view of the epistemic universe as being comprised of two inseparable entities, syntactic and semantic. Such a dual view of objects is well-established in mathematical logic where the syntactic notion of a formal theory is supplemented by the notion of a class of all its models. One could expect equally productive interactions between syntax and semantics in epistemology as well.

The definition of a game with epistemic conditions, cf. [6, 7], was originally semantic in a single-model format. In more recent papers (cf. [1, 9]), Aumann acknowledges the deficiencies of purely

semantic formalizations and asks for some kind of “syntactic epistemic logic” to bridge a gap between the syntactic character of game descriptions and the semantic way of analyzing games.

In this paper, we look at extensive games; the syntactic epistemic approach to strategic games has been tried in [2]. However, neither of these papers considers Epistemic Game Theory in its entirety, including probabilistic belief models, cf. [12]; we leave this for future studies.

3. Logical postulates and derivations

We consider the language of classical propositional logic augmented by modalities \mathbf{K}_i , for agent i 's knowledge, $i = 1, 2, \dots, n$. For the purposes of this paper, we consider the usual “knowledge postulates” (cf. [10, 13, 14, 16, 19]) corresponding to the multi-agent modal logic $S5_n$:⁴

- classical logic postulates and rule *Modus Ponens* $A, A \rightarrow B \vdash B$;
- distributivity: $\mathbf{K}_i(A \rightarrow B) \rightarrow (\mathbf{K}_i A \rightarrow \mathbf{K}_i B)$;
- reflection: $\mathbf{K}_i A \rightarrow A$;
- positive introspection: $\mathbf{K}_i A \rightarrow \mathbf{K}_i \mathbf{K}_i A$;
- negative introspection: $\neg \mathbf{K}_i A \rightarrow \mathbf{K}_i \neg \mathbf{K}_i A$;
- necessitation rule: $\vdash A \Rightarrow \vdash \mathbf{K}_i A$.

A derivation in $S5_n$ is a derivation from $S5_n$ -axioms by $S5_n$ -rules (*Modus Ponens* and necessitation). The notation

$$\vdash A \tag{4}$$

is used to represent the fact that A is derivable in $S5_n$.

3.1. Derivations from hypotheses

For a given set of formulas Γ (here called “hypotheses” or “assumptions”) we consider *derivations from* Γ : assume all $S5_n$ -theorems, Γ , and use classical reasoning (rule *Modus Ponens*). The notation

$$\Gamma \vdash A \tag{5}$$

represents A is derivable from Γ .

It is important to distinguish the role of necessitation in reasoning without assumptions (4) and in reasoning from a nonempty set of assumptions (5). In (4), necessitation can be used freely: what is derived from logical postulates ($\vdash A$) is known ($\vdash \mathbf{K}_i A$). In (5), the rule of necessitation is not postulated: if A follows from a set of assumptions Γ , we cannot conclude that A is known, since Γ itself can be unknown. However, for some “good” sets of assumptions Γ , necessitation is a valid rule (cf. Γ_3 from Example 4.2, MC_n from Section 5).

⁴The same approach works for other epistemic modal logics.

Example 3.1. If we want to describe a situation in which proposition m is known to agent 1, we consider the set of assumptions Γ :

$$\Gamma = \{\mathbf{K}_1 m\}.$$

From this Γ , by reflection principle $\mathbf{K}_1 m \rightarrow m$ from $S5_n$, we can derive m ,

$$\Gamma \vdash m.$$

Likewise, we can conclude ‘1 knows that 1 knows m ’ by using positive introspection:

$$\Gamma \vdash \mathbf{K}_1 \mathbf{K}_1 m.$$

However, we cannot conclude that agent 2 knows m :

$$\Gamma \not\vdash \mathbf{K}_2 m.$$

This is rather clear intuitively since when assuming ‘1 knows m ,’ we do not settle the question of whether ‘2 knows m .’⁵ Therefore, there is no necessitation in this Γ , since we have $\Gamma \vdash m$ but $\Gamma \not\vdash \mathbf{K}_2 m$.

3.2. Common knowledge and necessitation

We will also use abbreviations: for “everybody’s knowledge”

$$\mathbf{E}X = \mathbf{K}_1 X \wedge \dots \wedge \mathbf{K}_n X,$$

and “common knowledge”

$$\mathbf{C}X = \{X, \mathbf{E}X, \mathbf{E}^2 X, \mathbf{E}^3 X, \dots\}.$$

As one can see, $\mathbf{C}X$ is an infinite set of formulas. Since modalities \mathbf{K}_i commute with the conjunction, $\mathbf{C}X$ is provably equivalent to the set of all formulas which are X prefixed by iterated knowledge modalities:

$$\mathbf{C}X = \{P_1 P_2 \dots P_k X \mid k = 0, 1, 2, \dots, P_i \in \{\mathbf{K}_1, \dots, \mathbf{K}_n\}\}.$$

Naturally,

$$\mathbf{C}\Gamma = \bigcup \{\mathbf{C}F \mid F \in \Gamma\}$$

that states “ Γ is common knowledge.”

The set of formulas $\mathbf{C}X$ emulates common knowledge of X using the conventional modalities $\{\mathbf{K}_1, \dots, \mathbf{K}_n\}$. This allows us to speak, to the extent we need here, about common knowledge without introducing a special modality and new principles.

The following proposition states that the rule of necessitation corresponds to common knowledge of all assumptions. If Γ, Δ are sets of formulas, then $\Gamma \vdash \Delta$ means $\Gamma \vdash X$ for each $X \in \Delta$.

⁵A rigorous proof of this non-derivability can be made by providing a counter-model.

Proposition 3.2. *A set of formulas Γ is closed under necessitation if and only if $\Gamma \vdash \mathbf{C}\Gamma$, i.e., that Γ proves its own common knowledge.*

Proof:

Direction ‘if.’ Assume $\Gamma \vdash \mathbf{C}\Gamma$ and prove by induction on derivations that $\Gamma \vdash X$ yields $\Gamma \vdash \mathbf{K}_i X$. For X being a theorem of $\mathbf{S5}_n$, this follows from the rule of necessitation in $\mathbf{S5}_n$. For $X \in \Gamma$, it follows from the assumption that $\Gamma \vdash \mathbf{C}X$, hence $\Gamma \vdash \mathbf{K}_i X$. If X is obtained from *Modus Ponens*, $\Gamma \vdash Y \rightarrow X$ and $\Gamma \vdash Y$. By the induction hypothesis, $\Gamma \vdash \mathbf{K}_i(Y \rightarrow X)$ and $\Gamma \vdash \mathbf{K}_i Y$. By the distributivity principle of $\mathbf{S5}_n$, $\Gamma \vdash \mathbf{K}_i X$.

For ‘only if,’ suppose that Γ is closed under necessitation and $X \in \Gamma$, hence $\Gamma \vdash X$. Using appropriate instances of the necessitation rule in Γ we can derive $P_1 P_2 P_3, \dots, P_k X$ for each prefix $P_1 P_2 P_3, \dots, P_k$ with P_i is one of $\mathbf{K}_1, \mathbf{K}_2, \dots, \mathbf{K}_n$. Therefore, $\Gamma \vdash \mathbf{C}X$ and $\Gamma \vdash \mathbf{C}\Gamma$. \square

4. Kripke structures and models

A Kripke structure is a convenient vehicle for specifying epistemic assertions via truth values of atomic propositions and the combinatorial structure of the set of global states of the system. A **Kripke structure**

$$\mathcal{M} = \langle W, R_1, R_2, \dots, \Vdash \rangle$$

consists of a non-empty set W of possible worlds, “indistinguishability” equivalence relations R_1, R_2, \dots for each agent, and truth assignment ‘ \Vdash ’ of atoms at each world. The predicate ‘ F holds at u ’ ($u \Vdash F$) respects Booleans and reads epistemic assertions as

$$u \Vdash \mathbf{K}_i F \quad \text{iff} \quad \text{for each state } v \in W \text{ with } u R_i v, v \Vdash F \text{ holds.}$$

Conceptually, ‘agent i at state u knows F ’ ($u \Vdash \mathbf{K}_i F$) encodes the situation in which F holds at each state indistinguishable from u for agent i .

A **model** of a set of formulas Γ is a pair (\mathcal{M}, u) of a Kripke structure \mathcal{M} and a state u such that all formulas from Γ hold at u :

$$\mathcal{M}, u \Vdash F \text{ for all } F \in \Gamma.^6$$

A pair (\mathcal{M}, u) is an **exact model** of Γ if

$$\Gamma \vdash F \Leftrightarrow \mathcal{M}, u \Vdash F.$$

An epistemic scenario (a set of $\mathbf{S5}_n$ -formulas) Γ admits a **semantic definition** iff Γ has an exact model.

There is a simple criterion to determine whether Γ admits semantic definitions (Theorem 4.3) and we argue that “most” epistemic scenarios lack semantic definitions. These observations provide a justification for Syntactic Epistemic Logic with its syntactic approach to epistemic scenarios.

⁶Here we mean *local models* when Γ is satisfied at one world. This should not be confused with the more restrictive notion of *global models* when Γ is satisfied at each world of the model (cf. [10, 11, 15]). For epistemic purposes, global models do not suffice: for example, a consistent and meaningful situation $\Gamma = \{m, \neg \mathbf{K}m\}$, i.e., ‘ m holds but is not known’ does not have global models.

A formula F follows semantically from Γ ,

$$\Gamma \models F,$$

if F holds in each model (\mathcal{M}, u) of Γ . A well-known fact connecting syntactic derivability from Γ and semantic consequence is given by the **Completeness Theorem**⁷:

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

This fact has been used to claim the equivalence of the syntactic and semantic approaches and to define epistemic scenarios semantically by a model. However, the semantic part of the Completeness Theorem

$$\Gamma \models F$$

refers to the validity of F in **all** models of Γ , not in an arbitrary single model.

We challenge the model theoretical doctrine in epistemology and show the limitations of single-model semantic specifications, cf. Theorem 4.3.

4.1. Canonical model

The Completeness Theorem claims that if Γ does not derive F , then there is a model (\mathcal{M}, u) of Γ in which F is false. Where does this model come from?

The standard answer is given by the canonical model construction. In any model (\mathcal{M}, u) of Γ , the set of truths \mathcal{T} contains Γ and is *maximal*, i.e., for each formula F ,

$$F \in \mathcal{T} \quad \text{or} \quad \neg F \in \mathcal{T}.$$

This observation suggests the notion of a *possible world* as a maximal set of formulas Γ which is consistent, i.e., $\Gamma \not\vdash \perp$.

A *canonical model* $\mathcal{M}(S5_n)$ of $S5_n$ (cf. [10, 11, 13, 14, 15, 16]) consists of all possible worlds over $S5_n$. Accessibility relations are defined on the basis of what is known at each world: for maximal consistent sets α and β ,

$$\alpha R_i \beta \quad \text{iff} \quad \alpha_{\mathbf{K}_i} \subseteq \beta$$

where

$$\alpha_{\mathbf{K}_i} = \{F \mid \mathbf{K}_i F \in \alpha\},$$

i.e.,

$$\text{all facts that are known at } \alpha \text{ are true at } \beta.$$

Evaluations of atomic propositions are defined accordingly:

$$\alpha \Vdash p_i \quad \text{iff} \quad p_i \in \alpha.$$

⁷There are many sources in which the proof of this theorem can be found, e.g., [10, 11, 13, 14, 15, 16, 19].

The standard Truth Lemma shows that Kripkean truth values in the canonical model agree with possible worlds: for each formula F ,

$$\alpha \Vdash F \quad \text{iff} \quad F \in \alpha.$$

The canonical model $\mathcal{M}(S5_n)$ of $S5_n$ serves as a parametrized universal model for each consistent epistemic scenario Γ . Given Γ , by the well-known Lindenbaum construction, extend Γ to a maximal consistent set α . By definition, α is a possible world in $\mathcal{M}(S5_n)$. By the Truth Lemma, all formulas from Γ hold in α :

$$\mathcal{M}(S5_n), \alpha \Vdash \Gamma.$$

4.2. Deductive completeness

Definition 4.1. A set of $S5_n$ -formulas Γ is *deductively complete* if

$$\Gamma \vdash F \quad \text{or} \quad \Gamma \vdash \neg F.$$

Example 4.2. Consider examples in the language of the two-agent epistemic logic $S5_2$ with one propositional variable m and knowledge modalities \mathbf{K}_1 and \mathbf{K}_2 .

1. $\Gamma_1 = \{m\}$, where m is a propositional letter. Neither $\mathbf{K}_1 m$ nor $\neg \mathbf{K}_1 m$ is derivable in Γ_1 and this can be easily shown on corresponding models. Hence Γ_1 is not deductively complete.⁸

2. $\Gamma_2 = \{\mathbf{E}m\}$, i.e., both agents have first-order knowledge of m . However, the second-order knowledge assertions, e.g., $\mathbf{K}_2 \mathbf{K}_1 m$, are independent,⁹

$$\Gamma_2 \not\vdash \mathbf{K}_2 \mathbf{K}_1 m \quad \text{and} \quad \Gamma_2 \not\vdash \neg \mathbf{K}_2 \mathbf{K}_1 m.$$

This makes Γ_2 deductively incomplete.

3. $\Gamma_3 = \mathbf{C}m$, i.e., it is common knowledge that m . This set is deductively complete. Indeed, first note that, by Proposition 3.2, Γ_3 admits necessitation:¹⁰

$$\Gamma_3 \vdash F \quad \Rightarrow \quad \Gamma_3 \vdash \mathbf{K}_i F, \quad i = 1, 2.$$

To establish the completeness property: for each formula F ,

$$\Gamma_3 \vdash F \quad \text{or} \quad \Gamma_3 \vdash \neg F,$$

run induction on F . The base case when F is m is covered, since $\Gamma_3 \vdash m$. The Boolean cases are straightforward. Case $F = \mathbf{K}_i X$. If $\Gamma_3 \vdash X$, then, by necessitation, $\Gamma_3 \vdash \mathbf{K}_i X$. If $\Gamma_3 \vdash \neg X$, then, since $\mathbf{S5}$ proves $\neg X \rightarrow \neg \mathbf{K}_i X$, $\Gamma_3 \vdash \neg \mathbf{K}_i X$.

⁸In classical logic without epistemic modalities, Γ_1 is deductively complete: for each modal-free formula F of one variable m , either $\Gamma_1 \vdash F$ or $\Gamma_1 \vdash \neg F$.

⁹Again, there are easy countermodels.

¹⁰which is not the case for Γ_1 and Γ_2 .

4.3. Semantic definitions and complete scenarios

The following observation provides a necessary and sufficient condition for semantic definability. Let Γ be a consistent set of formulas in the language of $S5_n$.¹¹

Theorem 4.3. Γ is semantically definable if and only if it is deductively complete.

Proof:

The ‘only if’ direction. Suppose Γ has an exact model (\mathcal{M}, u) , i.e.,

$$\Gamma \vdash F \Leftrightarrow \mathcal{M}, u \Vdash F.$$

The set of true formulas in (\mathcal{M}, u) is maximal: for each formula F ,

$$\mathcal{M}, u \Vdash F \text{ or } \mathcal{M}, u \Vdash \neg F,$$

hence Γ is deductively complete: for each F ,

$$\Gamma \vdash F \text{ or } \Gamma \vdash \neg F.$$

The ‘if’ direction. Suppose Γ is consistent and deductively complete. Then the deductive closure $\tilde{\Gamma}$ of Γ

$$\tilde{\Gamma} = \{F \mid \Gamma \vdash F\},$$

is a maximal consistent set, hence an element of the canonical model $\mathcal{M}(S5_n)$. We claim that $(\mathcal{M}(S5_n), \tilde{\Gamma})$ is an exact model of Γ , i.e., for each F ,

$$\Gamma \vdash F \Leftrightarrow \mathcal{M}(S5_n), \tilde{\Gamma} \Vdash F.$$

Indeed, if $\Gamma \vdash F$, then $F \in \tilde{\Gamma}$ by the definition of $\tilde{\Gamma}$. By the Truth Lemma in $\mathcal{M}(S5_n)$, F holds at the world $\tilde{\Gamma}$. If $\Gamma \not\vdash F$, then, by deductive completeness of Γ , $\Gamma \vdash \neg F$, hence, as before, $\neg F$ holds at $\tilde{\Gamma}$, i.e., $\mathcal{M}(S5_n), \tilde{\Gamma} \not\Vdash F$. \square

Theorem 4.3 shows serious limitations of semantic definitions. Since, intuitively, deductively complete scenarios Γ are exceptions, “most” epistemic situations cannot be defined semantically.

In Section 5.4, we provide a yet another example of an incomplete but meaningful epistemic scenario, a natural variant of the Muddy Children puzzle, which, by Theorem 4.3 does not have a semantic definition, but can nevertheless be easily specified and analyzed syntactically.

In Section 6, we consider an example of an extensive game with incomplete epistemic description which cannot be defined semantically, but admits an easy syntactic analysis.

¹¹The same criteria hold for any other normal modal logic which has a canonical model in the usual sense.

5. The Muddy Children puzzle

Consider the standard Muddy Children puzzle, which is formulated syntactically.

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. The children see everybody else's foreheads, but not their own. The father says: "some of you are muddy," then adds: "Do any of you know that you have mud on your forehead? If you do, raise your hand now." No one raises a 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?

5.1. Standard syntactic formalization

This can be described in $S5_n$ with modalities $\mathbf{K}_1, \mathbf{K}_2, \dots, \mathbf{K}_n$ for the children's knowledge and atomic propositions m_1, m_2, \dots, m_n with m_i stating "child i is muddy." The **initial configuration**, which we call MC_n , includes common knowledge assertions of the following 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 themselves:*

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

Transition from the verbal description of the situation to MC_n is a straightforward formalization of a given syntactic description to another, logic friendly syntactic form.

5.2. Semantic solution

In the standard semantic solution, the set of assumptions MC_n is replaced by a Kripke model: n -dimensional cube Q_n ([14, 16, 17, 18, 19]). To keep things simple, we consider the case $n = k = 2$.




Figure 2. Model Q_2 .

Logical possibilities for the truth value combinations¹² of (m_1, m_2) , namely $(0,0)$, $(0,1)$, $(1,0)$, and $(1,1)$ are declared possible worlds. There are two indistinguishability relations denoted by solid arrows (for 1) and dotted arrows (for 2). It is easy to check that conditions 1 (knowing about the others) and 2 (not knowing about themselves) hold at each node of this model. Furthermore, Q_2 is assumed to be commonly known.




Figure 3. Models \mathcal{M}_2 and \mathcal{M}_3 .

After the father publicly announces $m_1 \vee m_2$, node $(0,0)$ is no longer possible and model \mathcal{M}_2 now becomes common knowledge. Both children realize that in $(1,0)$, child 2 would know whether (s)he is muddy (no other 2-indistinguishable worlds), and in $(0,1)$, child 1 would know whether (s)he is muddy. After both children answer “no” to whether they know what is on their foreheads, worlds $(1,0)$ and $(0,1)$ are no longer possible, and each child eliminates them. The only remaining logical possibility here is model \mathcal{M}_3 . Now both children know that their foreheads are muddy.

5.3. Justifying the model

The semantic solution in Section 5.2 adopts Q_n as a semantic equivalent of a theory MC_n . Can this choice of the model be justified? In the case of Muddy Children, the answer is ‘yes.’

Let u be a node at Q_n . i.e., u is an n -tuple of 0’s and 1’s and $u \Vdash m_i$ iff i ’s projection of u is 1. Naturally, u is represented by a formula $\pi(u)$:

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

It is obvious that $v \Vdash \pi(u)$ iff $v = u$.

By $MC_n(u)$ we understand the Muddy Children scenario with specific distribution of truth values of m_i ’s corresponding to u :

$$MC_n(u) = MC_n \cup \{\pi(u)\}.$$

So, each specific instance of Muddy Children is formalized by an appropriate $MC_n(u)$.

Theorem 5.1. *Each instance $MC_n(u)$ of Muddy Children is deductively complete and (Q_n, u) is its exact model*

$$MC_n(u) \vdash F \quad \text{iff} \quad Q_n, u \Vdash F.$$

¹²1 standing for ‘true’ and 0 for ‘false’

Proof:¹³ The direction ‘only if’ claims that (Q_n, u) is a model for $\text{MC}_n(u)$ is straightforward. First, Q_n is an S5_n -model and all principles of S5_n hold everywhere in Q_n . It is easy to see that principles ‘knowing about the others’ and ‘not knowing about himself’ hold at each node. Furthermore, as $\pi(u)$ holds at u , everything that can be derived from $\text{MC}_n(u)$ holds at u .

To establish the ‘if’ direction, we first note that, by Proposition 3.2, necessitation is admissible in MC_n : for each F ,

$$\text{MC}_n \vdash F \Rightarrow \text{MC}_n \vdash \mathbf{K}_i F.$$

The theorem now follows from the statement $\mathcal{S}(F)$:

for all nodes $u \in Q_n$,

$$Q_n, u \Vdash F \Rightarrow \text{MC}_n \vdash \pi(u) \rightarrow F$$

and

$$Q_n, u \Vdash \neg F \Rightarrow \text{MC}_n \vdash \pi(u) \rightarrow \neg F.$$

We prove that $\mathcal{S}(F)$ holds for all F by induction on F .

The case F is one of the atomic propositions m_1, m_2, \dots, m_n is trivial since $\text{MC}_n \vdash \pi(u) \rightarrow m_i$, if $u \Vdash m_i$ and $\text{MC}_n \vdash \pi(u) \rightarrow \neg m_i$, if $u \Vdash \neg m_i$. The Boolean cases are also straightforward.

The case $F = \mathbf{K}_i X$. Consider the node u^i which differs from u only at the i -coordinate. Without a loss of generality, we may assume that $u \Vdash m_i$ and $u^i \Vdash \neg m_i$; the alternative $u \Vdash \neg m_i$ and $u^i \Vdash m_i$ is similar.

Suppose $Q_n, u \Vdash \mathbf{K}_i X$. Then $Q_n, u \Vdash X$ and $Q_n, u^i \Vdash X$. By the induction hypothesis,

$$\text{MC}_n \vdash \pi(u) \rightarrow X \quad \text{and} \quad \text{MC}_n \vdash \pi(u^i) \rightarrow X.$$

By the rules of logic (splitting premises)

$$\text{MC}_n \vdash \pi(u)_{-i} \rightarrow (m_i \rightarrow X) \quad \text{and} \quad \text{MC}_n \vdash \pi(u)_{-i} \rightarrow (\neg m_i \rightarrow X),$$

where $\pi(v)_{-i}$ is $\pi(v)$ without its i -th coordinate¹⁴. By further reasoning,

$$\text{MC}_n \vdash \pi(u)_{-i} \rightarrow X.$$

By necessitation in MC_n , and distributivity,

$$\text{MC}_n \vdash \mathbf{K}_i \pi(u)_{-i} \rightarrow \mathbf{K}_i X.$$

By ‘knowing about the others’ principle, and since $\pi(u)_{-i}$ contains only atoms other than m_i ,

$$\text{MC}_n \vdash \pi(u)_{-i} \rightarrow \mathbf{K}_i \pi(u)_{-i},$$

¹³We have chosen to present a syntactic proof of Theorem 5.1. A semantic proof that makes use of bi-simulations can also be given.

¹⁴Formally, $\pi(v)_{-i} = \bigwedge \{m_j \mid v \Vdash m_j, j \neq i\} \wedge \bigwedge \{\neg m_j \mid v \Vdash \neg m_j, j \neq i\}$.

hence

$$\text{MC}_n \vdash \pi(u)_{-i} \rightarrow \mathbf{K}_i X,$$

and

$$\text{MC}_n \vdash \pi(u) \rightarrow \mathbf{K}_i X.$$

Now suppose $Q_n, u \Vdash \neg \mathbf{K}_i X$. Then $Q_n, u \Vdash \neg X$ or $Q_n, u^i \Vdash \neg X$. By the induction hypothesis,

$$\text{MC}_n \vdash \pi(u) \rightarrow \neg X \quad \text{or} \quad \text{MC}_n \vdash \pi(u^i) \rightarrow \neg X.$$

In the former case we immediately get $\text{MC}_n \vdash \pi(u) \rightarrow \neg \mathbf{K}_i X$, by reflection $\neg X \rightarrow \neg \mathbf{K}_i X$. So, consider the latter, i.e., $\text{MC}_n \vdash \pi(u^i) \rightarrow \neg X$. As before,

$$\text{MC}_n \vdash \pi(u)_{-i} \rightarrow (\neg m_i \rightarrow \neg X).$$

By contrapositive,

$$\text{MC}_n \vdash \pi(u)_{-i} \rightarrow (X \rightarrow m_i).$$

By necessitation and distribution,

$$\text{MC}_n \vdash \mathbf{K}_i \pi(u)_{-i} \rightarrow (\mathbf{K}_i X \rightarrow \mathbf{K}_i m_i).$$

By ‘knowing about others,’ as before,

$$\text{MC}_n \vdash \pi(u)_{-i} \rightarrow (\mathbf{K}_i X \rightarrow \mathbf{K}_i m_i).$$

By ‘not knowing about himself,’ $\text{MC}_n \vdash \neg \mathbf{K}_i m_i$, hence

$$\text{MC}_n \vdash \pi(u)_{-i} \rightarrow \neg \mathbf{K}_i X,$$

and

$$\text{MC}_n \vdash \pi(u) \rightarrow \neg \mathbf{K}_i X.$$

As we see, in the case of Muddy Children given by a syntactic description, $\text{MC}_n(u)$, picking one “natural model” (Q_n, u) could be justified. However, in a general setting, the approach

given a syntactic description, pick a “natural model”

is intrinsically flawed: by Theorem 4.3, in many (intuitively, most) cases, there is no model description at all. Furthermore, if there is a “natural model,” a completeness analysis in the style of what we did for MC_n in Theorem 5.1 is required.

5.4. Incomplete scenario: Muddy Children Explicit

Here is a natural modification, $MCE_{n,k}$, of the standard Muddy Children.

A group of n children meet their father after playing in the mud. Each child sees everybody else's foreheads. The father says: "k of you are muddy" after which it becomes common knowledge that all children know whether they are muddy. Why?

This description does not specify whether children initially know if they are muddy; hence the initial description of $MCE_{n,k}$ is, generally speaking, not complete¹⁵. By Theorem 4.3, the initial $MCE_{2,2}$ is not semantically definable. Therefore, $MCE_{2,2}$ cannot be treated by "natural model" methods.

However, here is a syntactic analysis of $MCE_{n,k}$ which can be shaped as a formal logical reasoning within an appropriate extension of $S5_n$.

After father's announcement, each child knows that if she sees k muddy foreheads, then she is not muddy, and if she sees $k - 1$ muddy foreheads, she is muddy: this secures that each child knows whether she is muddy. Moreover, everybody can reflect on this reasoning and this makes it common knowledge that each child knows whether she is muddy.

5.5. Some additional observations

If we want to go beyond complete epistemic scenarios, we need a mathematical apparatus to handle classes of models, and not just single models. The format of syntactic specifications in some version of the modal epistemic language is a viable candidate for such an apparatus.

The traditional model solution of MC_n without completeness analysis uses a strong additional assumption – common knowledge of a specific model Q_n and hence, strictly speaking, does not resolve the original Muddy Children puzzle; it rather corresponds to a different scenario of a more tightly controlled epistemic states of agents, e.g.,

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

One could argue that the given model solution of MC_n actually codifies some 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.

6. Syntactic Epistemic Logic and games

Consider a variant *Centipede Lite*, CL, of the well-known Centipede game (cf. [17]) with risk-averse rational players Alice and Bob. No cross-knowledge of rationality, let alone common knowledge, is assumed!

¹⁵In particular, prior to father's announcement $MCE_{2,2}$ does not specify whether $K_1 m_1$ holds or not.




Figure 4. Centipede game tree

CL admits the following rigorous analysis.

At 3, Alice plays down. At 2, Bob plays down because he is risk-averse and cannot rule out that Alice plays down at 3 (since it is true). At 1, Alice plays down because she cannot rule out Bob's playing down at 2. So, CL has the so-called Backward Induction solution “*down at each node.*”

CL is not complete (epistemic assumptions, such as *Bob knows that Alice plays “across” at 3*, are not specified), hence CL cannot be defined by a single Kripke/Aumann model.

7. Incomplete and complete scenarios

How typical are deductively incomplete epistemic scenarios? We argue that this is the rule rather than the exception. Epistemic conditions more flexible than common knowledge of the game and rationality (mutual knowledge of rationality, asymmetric epistemic assumptions when one player knows more than the other, etc.) lead to semantic undefinability.

Semantically non-definable scenarios are the “dark matter” of the epistemic universe: they are everywhere, but cannot be visualized as a single model. The semantic approach does not recognize these “dark matter” scenarios; SEL deals with them syntactically.

The question remains: how manageable are semantic definitions of deductively complete scenarios?

7.1. Cardinality and knowability issue

Models of complete Γ 's provided by Theorem 4.3 are instances of the canonical model $\mathcal{M}(S5_n)$ at nodes $\tilde{\Gamma}$ corresponding to Γ . This generic solution is, however, not satisfactory because of the highly nonconstructive nature of the canonical model $\mathcal{M}(S5_n)$.

As was shown in [8], the canonical model $\mathcal{M}(S5_n)$ for any $n \geq 1$ has continuum-many possible worlds even with just one propositional letter. This alone renders models $(\mathcal{M}(S5_n), \tilde{\Gamma})$ not knowable under any reasonable meaning of “known.” The canonical model for $S5_n$ is just too large to be considered known and hence does not *a priori* satisfy the knowability of the model requirement II from Section 1.

This observation suggests that the question about existence of an epistemically acceptable (“known”) model for a given deductively complete set Γ requires a case-by-case consideration.

7.2. Complexity considerations

Epistemic models of even simple and complete scenarios can be prohibitively large compared to their syntactic descriptions. For example, the Muddy Children model Q_n is exponential in n whereas its syntactic description MC_n is quadratic in n .

Consider a real-life epistemic situation after the cards have been initially dealt in the game of poker. One can show that for each distribution of cards, its natural syntactic description in epistemic logic is deductively complete ([5]) and hence admits a model characterization. Moreover, it has a natural finite model of the type given in [14] with hands as possible worlds and with straightforward knowledge relations. However, with 52 cards and 4 players there are over 10^{24} different combinations of hands. This yields that explicit formalization of the model not practical. Players reason using concise syntactic descriptions of the rules of poker and of its “large” model in the natural language, which can also be syntactically formalized in some kind of extension of epistemic logic.

In this and some other real life situations, models are prohibitively large whereas appropriate syntactic descriptions can be quite manageable.

8. Further observations

An interesting question is why the traditional semantic approach, despite its aforementioned shortcomings, produces correct answers in many situations. One of possible reasons for this is *pragmatic self-limitation*.

Given a syntactic description \mathcal{D} , we intuitively seek a solution that logically follows from \mathcal{D} . Even if we reason on a “natural model” of \mathcal{D} , normally overspecified, we try not to use features of the model that are not supported by \mathcal{D} . If we conclude a property P by such self-restricted reasoning about the model, then P indeed logically follows from \mathcal{D} .

This situation resembles Geometry, in which we reason about “models”, i.e., combinations of triangles, circles, etc., but have a rigorous system of postulates in the background. We are trained not to venture beyond given postulates even in informal reasoning.

Such an *ad hoc* pragmatic approach needs a scientific foundation, which could be provided within the framework of Syntactic Epistemic Logic.

9. Syntactic Epistemic Logic suggestions

The Syntactic Epistemic Logic suggestion, in brief, is *to make an appropriate syntactic formalization of an epistemic scenario its formal specification*. This extends the scope of scientific epistemology and offers a remedy for two principal weaknesses of the traditional semantic approach. The reader will recall that those weaknesses were the restricting single model requirement and a hidden assumption of the common knowledge of this model.

SEL suggests a way to handle incomplete scenarios which have rigorous syntactic descriptions (cf. Muddy Children Explicit, Centipede Lite, etc.).

SEL offers a scientific framework for resolving the tension, identified by R. Aumann [9], between a syntactic description and its hand-picked model. If, given a syntactic description Γ we prefer to reason on a model \mathcal{M} , we have to establish completeness of Γ with respect to \mathcal{M} .

Appropriate syntactic specifications could also help to handle situations for which natural models exist but are too large for explicit presentations.

SEL can help to extend Epistemic Game Theory to less restrictive epistemic conditions. A broad class of epistemic scenarios does not define higher-order epistemic assertions and rather addresses individual knowledge, mutual and limited-depth knowledge, asymmetric knowledge, etc. and hence is deductively incomplete and has no exact single model characterizations. However, if such a scenario allows a syntactic formulation, it can be handled scientifically by a variety of mathematical tools, including reasoning about its models.

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

Acknowledgements

The author is grateful to Adam Brandenburger, Alexandru Baltag, Johan van Benthem, Robert Constable, Melvin Fitting, Vladimir Krupski, Anil Nerode, Elena Nogina, Eoin Moore, Vincent Peluce, Tudor Protopopescu, Bryan Renne, Richard Shore, and Cagil Tasdemir for useful discussions. Special thanks to Karen Kletter for editing early versions of this text.

References

- [1] Arieli I, Aumann R. The logic of backward induction. doi:10.2139/ssrn.2133302, 2012.
- [2] Artemov S. On Definitive Solutions of Strategic Games. Alexandru Baltag and Sonja Smets, eds. *Johan van Benthem on Logic and Information Dynamics*. Outstanding Contributions to Logic 5:487–507, Springer, 2014. doi: 10.1007/978-3-319-06025-5_17.
- [3] Artemov S. Syntactic Epistemic Logic. *Book of Abstracts. 15th Congress of Logic, Methodology and Philosophy of Science, University of Helsinki, 2015*, 109–110.
- [4] Artemov S. Hyperderivations. *The Hausdorff Trimester Program: Types, Sets and Constructions, Hausdorff Center for Mathematics, Bonn, 2018*.
<https://www.youtube.com/watch?v=kytYAi6Ln7U&t=1276s>
- [5] Artemov S, Nogina E. On completeness of epistemic theories. *The Bulletin of Symbolic Logic*, 24(2):232, 2018. <https://doi.org/10.1017/bsl.2018.13>.
- [6] Aumann R. Agreeing to disagree. *The Annals of Statistics*, 4(6):1236–1239, 1976. <https://doi.org/10.1214/aos/1176343654>.
- [7] Aumann R. Backward induction and common knowledge of rationality. *Games and Economic Behavior*, 8(1):6–19, 1995. [https://doi.org/10.1016/S0899-8256\(05\)80015-6](https://doi.org/10.1016/S0899-8256(05)80015-6).
- [8] Aumann R. Interactive epistemology I: Knowledge. *International Journal of Game Theory*, 28:263–300, 1999. <https://doi.org/10.1007/s001820050111>.

- [9] Aumann R. Epistemic Logic: 5 Questions, 2010. Vincent F. Hendricks and Olivier Roy, eds. Automatic Press/VIP, pp. 21–33. ISBN 8792130240, 9788792130242
- [10] Blackburn P, de Rijke M, Venema Y. Modal Logic. *Cambridge Tracts in Theoretical Computer Science*, **53**, 2001.
- [11] Blackburn P, van Benthem J. Modal logic: A semantic perspective. *Handbook of Modal Logic*. pp.1–84. *Studies in Logic and Practical Reasoning* **3**, Elsevier, 2007. [https://doi.org/10.1016/S1570-2464\(07\)80004-8](https://doi.org/10.1016/S1570-2464(07)80004-8)
- [12] Brandenburger A. The Language of Game Theory: Putting Epistemics Into the Mathematics of Games. World Scientific Publishing Company, 2014. ISSN 2251-2071.
- [13] Chagrov A, Zakharyashev M. Modal Logic. *Oxford Logic Guides* **35**, 1997. ISBN-13: 978-0198537793; ISBN-10: 0198537794.
- [14] Fagin R, Halpern J, Moses Y, Vardi M. Reasoning About Knowledge. MIT Press, 1995. ISBN-13: 978-0262562003; ISBN-10: 9780262562003.
- [15] Fitting M. Modal proof theory. *Handbook of Modal Logic*. pp.85–138. *Studies in Logic and Practical Reasoning* **3**, Elsevier, 2007. [https://doi.org/10.1016/S1570-2464\(07\)80005-X](https://doi.org/10.1016/S1570-2464(07)80005-X)
- [16] Meyer JJC, van der Hoek W. Epistemic Logic for AI and Computer Science. *Cambridge Tracts in Theoretical Computer Science* **41**, 1995.
- [17] Osborne M, Rubinstein A. A Course in Game Theory. MIT Press, 1994.
- [18] Pauly M, van der Hoek W. Modal logic for games and information. *Handbook of Modal Logic*. pp.1077–1148. *Studies in Logic and Practical Reasoning* **3**, Elsevier, 2007. [https://doi.org/10.1016/S1570-2464\(07\)80023-1](https://doi.org/10.1016/S1570-2464(07)80023-1)
- [19] Van Benthem J. Logic in Games. MIT Press, 2014.