Syntactic Epistemic Logic and Games

Sergei Artemov

The Graduate Center, The City University of New York
365 Fifth Avenue, New York City, NY 10016
sartemov@gc.cuny.edu

April 15, 2016

Abstract

The traditional representation of an epistemic scenario as a model covers only complete descriptions that specify truth values of all assertions. However, many, perhaps most, epistemic scenarios are not complete and allow partial or asymmetric knowledge. Syntactic Epistemic Logic, SEL, suggests viewing an epistemic situation as a set of syntactic conditions rather than as a model, thus also capturing incomplete descriptions. This helps to extend the scope of Epistemic Game Theory. In addition, SEL closes the conceptual and technical gap, identified by R. Aumann, between the syntactic character of game descriptions and the semantic method of analyzing games.

1 Introduction

We intend to promote changing the way logic and epistemic-related applications, in particular, in game theory, specify epistemic scenarios. We argue that the traditional semantic specifications as a single Kripke or Aumann structure are too restrictive since they cover only deductively complete descriptions, demanding *de facto* specification of truth values of all sentences in the language, not just the relevant ones.

Furthermore, in a normal situation in which a scenario is originally described syntactically and then formalized as a model, a completeness analysis relating these two modes of formalization is required but is basically never done, which is one of the shortcomings of the semantic approach.

For example, the Muddy Children puzzle is given syntactically but the reasoning is performed on a specific model, presumed commonly known, without rigorous justification (cf. [15, 19, 20, 21, 22]). We will show, that the choice of this model for Muddy Children can be justified, cf. Section 4. However, this case is a fortuitous exception: see Sections 4.6, 5.1 for syntactically defined scenarios without exact semantic definitions.

1.1 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 semantical 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 given $\mathcal{I}$, proceed to a specific epistemic model $\mathcal{M}_{\mathcal{I}}$, and make the latter a mathematical definition of $\mathcal{I}$:

$$\textit{informal description } \mathcal{I} \;\;\Rightarrow\;\; \textit{“natural model” } \mathcal{M}_{\mathcal{I}}. \tag{1}$$

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. Epistemic scenarios outside this group, which include situations with asymmetric and less-than-common knowledge (e.g., mutual knowledge) of conditions, typically do not have acceptable single-model presentations, but can be specified and handled syntactically.

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** $\mathcal{I}$:

$$\textit{description } \mathcal{I} \;\;\Rightarrow\;\; \textit{syntactic formalization } \mathcal{S}_{\mathcal{I}} \Rightarrow \textit{all its models } \mathcal{M}_{\mathcal{S}}. \tag{2}$$

The first step from $\mathcal{I}$ to $\mathcal{S}_{\mathcal{I}}$ is a formalization problem which certainly has its own subtleties which we, however, will not analyze here[1]. The SEL approach

[1]except for a brief discussion in Section 4.7 about representing negative knowledge.

(2), we argue, encompasses a broader class of epistemic scenarios than a semantic approach (1).

1.2 From semantic to syntactic in game theory

The rigorous definition of a game with epistemic knowledge conditions, cf. [6, 7], was originally purely semantic in a single-model format. Later, in [8], Aumann discusses two modes of formalization of epistemic conditions, semantic and syntactic, and argues that they are basically equivalent. 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.

The current paper contributes to this development by outlining, from the position of logic, the Syntactic Epistemic Logic framework. In addition to sealing the gap between syntactic and semantic, SEL extends the scope of Epistemic Game Theory by capturing natural scenarios which are not represented semantically by single models.

Furthermore, SEL does not lose efficient semantical descriptions. Each model yields the set of valid formulas which may be regarded as a complete syntactic description $\mathcal{T}$ of the situation. Moreover, if the initial model is finite, then $\mathcal{T}$ is linear-time decidable (cf. Section 7.2). So, efficient semantic descriptions are also accepted in SEL.

If a scenario is originally described syntactically as a set of assumptions Γ but we wish to work with a specific model $\mathcal{M}$, then we should establish the completeness of Γ with respect to $\mathcal{M}$, cf. the Muddy Children example in Section 4.

1.3 Is the syntactic approach practical?

SEL provides a balanced view of the epistemic universe as being comprised of two inseparable entities, syntactic and semantic. Such a dual picture is well-tested and has been highly beneficial in many contexts such as mathematics, logic, linguistics, etc. The 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 model of a theory, with rich mutually beneficial connections. One could expect equally productive interactions between syntax and semantics in epistemology as well.

1.4 Disclaimer

In this paper, we apply SEL to extensive games; the syntactic epistemic approach to strategic games has been outlined in [4]. However, neither of these papers considers Epistemic Game Theory in its entirety, including probabilistic belief models, cf. [13]; we leave this for future studies.

We realize that the list of references is not adequate to the task of paradigm change and does not do justice to the volumes of quality work in Epistemic Logic and Epistemic Game Theory. We are determined to improve this in the future.

2 Logical postulates and derivations

We consider the language of classical logic[2], augmented by modalities $\mathbf{K}_i$, for agent i's knowledge, $i = 1, 2, \ldots, n$. For the purposes of this paper, we assume the usual "knowledge postulates" (cf. [11, 14, 15, 19, 22]) corresponding to the multi-agent modal logic $\mathsf{S5}_n$:

- 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: $\vdash A \quad \Rightarrow \quad \vdash \mathbf{K}_i A$.

Derivation in $\mathsf{S5}_n$ is a derivation from $\mathsf{S5}_n$-axioms by $\mathsf{S5}_n$-rules (*modus ponens* and necessitation). This suffices to formalize any rigorous reasoning from the above principles. The notation

$$\vdash A \tag{3}$$

is used to represent the fact that A is derivable in $\mathsf{S5}_n$. The necessitation rule reflects the principle that all logical truths, i.e., $\mathsf{S5}_n$-principles and hence all their consequences ($\mathsf{S5}_n$-theorems) are common knowledge.

[2]In this paper, we consider the propositional case: symbol $\perp$ stands for the proposition 'false.'

An easy but useful observation: knowledge modalities and conjunction commute:

$$\vdash \mathbf{K}_i(A \wedge B) \leftrightarrow (\mathbf{K}_i A \wedge \mathbf{K}_i B).$$

2.1 Derivations from hypotheses

For a given set of formulas Γ (here called "hypotheses" or "assumptions") we consider *derivations from* Γ: assume all logical truths (i.e., the set of $\mathsf{S5}_n$-theorems) and Γ and use classical reasoning (rule *modus ponens*). The notation

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

represents 'A is derivable from Γ.'

It is important to distinguish the role of necessitation in reasoning without assumptions (3) and in reasoning from a nonempty set of assumptions (4). In (3), necessitation can be used freely: what is derived from logical postulates ($\vdash A$) is known ($\vdash \mathbf{K}_i A$). In (4), 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 2, MC_n from Section 4, Game 1 from Section 5).

Example 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 $\mathsf{S5}_n$, 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.[3] Therefore, there is no necessitation in this Γ, since we have $\Gamma \vdash m$ but $\Gamma \not\vdash \mathbf{K}_2 m$.

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

2.2 Common knowledge and necessitation

We will also use abbreviations: for "everybody's knowledge"

$$\mathbf{E}X = \mathbf{K}_1 X \wedge \mathbf{K}_2 X,$$

and "common knowledge"

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

As one can see, $\mathbf{C}X$ is an infinite (though quite regular and decidable) set of formulas. Since modalities $\mathbf{K}_i$ commute with the conjunction $\wedge$, $\mathbf{C}X$ is the set of all formulas which are X prefixed by iterated knowledge modalities:

$$\mathbf{C}X = \{P_1 P_2 \ldots P_k X \mid k = 0, 1, 2, \ldots,\quad P_i \in \{\mathbf{K}_1, \mathbf{K}_2\}\}.$$

Naturally, $\mathbf{C}\Gamma = \bigcup\{\mathbf{C}F \mid F \in \Gamma\}$ that states "Γ is common knowledge." The following proposition states that the rule of necessitation corresponds to common knowledge of all derivable facts.

Proposition 1 *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 from $\mathsf{S5}_n$, this follows from the rule of necessitation in $\mathsf{S5}_n$. 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 $\mathsf{S5}_n$, $\Gamma \vdash \mathbf{K}_i X$.

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

3 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.

Kripke structure[4]

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

consists of a non-empty set W of possible worlds, "indistinguishability" equivalence relations $R_1, R_2, \ldots$ for each agent, and truth assignment '$\Vdash$' of atoms at each world. Predicate '*F holds at u*' ($u \Vdash F$) respects Booleans and reads epistemic assertions as

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

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:

$$\mathcal{M}, u \Vdash F \ \ \textit{for all } F \in \Gamma.$$[5]

The **soundness property** states that if $(\mathcal{M}, u)$ is a model of Γ, then

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

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

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

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

There is a simple criterion to determine whether Γ admits semantical definitions (Theorem 1) and we argue that "most" epistemic scenarios lack semantical definitions. These observations provide a justification for Syntactic Epistemic Logic with its syntactic definitions of epistemic scenarios.

A formula F follows semantically from Γ,

$$\Gamma \models F,$$

[4]Introduced for multi-agent case in [17].

[5]Here 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. [11, 12, 16]). 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.

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**[6]:

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

This fact has been used by some to claim the equivalence of the syntactic and semantic approaches and to define epistemic scenarios semantically by a model. We challenge these claims and show the limitations of single-model semantic specifications, cf. Theorem 1.

3.1 Aumann knowledge structures

A *knowledge structure* (a.k.a. *partition structure*) is a tuple

$$\mathcal{M} = (\Omega, \mathcal{K}_1, \ldots, \mathcal{K}_n, \mathbf{s}),$$

where Ω is a set of "global states," $\mathcal{K}_1, \ldots, \mathcal{K}_n$ are "knowledge partitions" of Ω corresponding to players $1, 2, \ldots, n$, and $\mathbf{s}$ is a mapping from Ω to the set of all strategy profiles[7]: for a state ω,

$$\mathbf{s}(\omega) = (s_1, \ldots, s_n).$$

If we associate propositional variables with strategies, the strategy profile function $\mathbf{s}$ assigns truth values to these variables in the worlds of the model. So, the truth of atomic propositions is defined at each state $\omega \in \Omega$, the Boolean connectives behave conventionally, and

$$\mathbf{K}_i F \textit{ holds at } \omega \quad \text{iff} \quad F \textit{ holds at any } \omega' \in \mathcal{K}_i(\omega),$$

i.e., F holds at any ω' indistinguishable from ω for player i.

With these assumptions, and equivalence relations instead of partitions, knowledge structures may be regarded as Kripke $\mathsf{S5}$-models.

The fundamental requirement of knowledge structures (as well as Kripke models) in epistemic contexts is common knowledge of the model [8, 9][8]. For example, to interpret

$$u \Vdash \mathbf{K}_2 \mathbf{K}_1 p$$

[6] There are many sources in which the proof of this theorem can be found, e.g., [11, 12, 14, 15, 16, 19, 22].

[7] Cf. Section 5 for more details about strategies and profiles.

[8] Knowledge of the "real" world is not required.

as

at world u agent 2 knows that agent 1 knows p,

we have to assume that

> *when at u, 2 knows that at each v which is 2-equivalent to u,* 1 *knows that for each w which is 1-equivalent to v, p holds.*

This requires that 2 knows both the model and that 1 knows the model, etc.

Common knowledge of the model is naturally required when the model is a vehicle of defining epistemic values at each world:

$$\textit{model} \quad \Rightarrow \quad \textit{epistemic assertions at each world.}$$

3.2 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 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., $\Gamma \not\vdash \bot$.

A *canonical model* $\mathcal{M}(\mathsf{S5}_n)$ of $\mathsf{S5}_n$ (cf. [11, 12, 14, 15, 16, 19]) consists of all possible worlds over $\mathsf{S5}_n$. We will also be interested in versions of $\mathsf{S5}_n$ with finite sets of propositional letters, e.g., $\{m_1, m_2, \ldots, m_k\}$. Accessibility relations are defined on the basis of what is known at each world: for maximal consistent 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.,

all facts that are i-known at α are true at β.

The standard and well-known reasoning demonstrates that, since all theorems of $\mathsf{S5}_n$ belong to each possible world, each R_i is an equivalence relation. Evaluations of atomic propositions are defined according to the possible worlds:

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

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

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

Canonical model $\mathcal{M}(\mathsf{S5}_n)$ serves as a universal model of all epistemic scenarios. The following proposition is a well-known model existence property of syntactic descriptions.

Proposition 2 *Any consistent set of formulas* Γ *has a model.*

Proof. Given consistent Γ, by the well-known Lindenbaum construction, extend Γ to a maximal consistent set α. By definition, α is a possible world in $\mathcal{M}(\mathsf{S5}_n)$. By the Truth Lemma, all formulas from Γ hold in α:

$$\mathcal{M}(\mathsf{S5}_n), \alpha \Vdash \Gamma.$$

$\square$

We see that the canonical model $\mathcal{M}(\mathsf{S5}_n)$ of $\mathsf{S5}_n$ plays a crucial role here: it provides a Kripke model for each consistent set of formulas Γ with an appropriate choice of α.

However, $\mathcal{M}(\mathsf{S5}_n)$ itself is not an independently defined semantical structure for specifying knowledge assertions. On the contrary, states and relations of the canonical model are reverse engineered from syntactic data of what is assumed to be "known" at each possible world:

$$\textit{truth values of epistemic assertions} \quad \Rightarrow \quad \textit{canonical model.}$$

Possible world α from proof of Proposition 2 is also built syntactically (the Lindenbaum construction) from a given Γ and the resulting model $(\mathcal{M}(\mathsf{S5}_n), \alpha)$ is defined given Γ and hence is not a semantic definition of Γ.

Another feature of $\mathcal{M}(\mathsf{S5}_n)$ is that this model is not constructively defined and has continuum-many states even in the simplest case with one propositional letter and $n > 1$ [8]. As a result, $\mathcal{M}(\mathsf{S5}_n)$ cannot be considered known as a Kripke/Aumann structure in any reasonable sense of "known."

3.3 Deductive completeness

Definition 1 A set of $\mathsf{S5}_n$-formulas Γ is *deductively complete* if

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

Example 2 Consider examples in the language of the two-agent epistemic logic $\mathsf{S5}_2$ 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[9]. Hence Γ_1 is not deductively complete.

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[10]

$$\Gamma_2 \nvdash \mathbf{K}_2\mathbf{K}_1 m \quad \text{and} \quad \Gamma_2 \nvdash \neg \mathbf{K}_2\mathbf{K}_1 m.$$

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 for formulas over variable m. Indeed, first note that, by Proposition 1, Γ_3 admits necessitation[11]:

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

To establish the completeness property: for each formula F,

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

run an 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 $\mathsf{S5}$ proves $\neg X \rightarrow \neg \mathbf{K}_i X$, $\Gamma_3 \vdash \neg \mathbf{K}_i X$.

[9] Note that in classical logic without epistemic modalities, this set Γ_1 is deductively complete: for each modal-free formula F of one variable m, either $\Gamma_1 \vdash F$ or $\Gamma_1 \vdash \neg F$. So, epistemic modalities make the difference here.

[10] Again, there are easy countermodels.

[11] which is not the case for Γ_1 and Γ_2.

3.4 Semantical definitions and complete scenarios

The following theorem provides a necessary and sufficient condition for definability. Let Γ be a consistent set of formulas in the language of epistemic logic $\mathsf{S5}_n$.[12]

Theorem 1 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.,

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

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

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

hence Γ is deductively complete: for each F,

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

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

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

is a maximal consistent set, hence an element of the canonical model $\mathcal{M}(\mathsf{S5}_n)$. We claim that $(\mathcal{M}(\mathsf{S5}_n), \widetilde{\Gamma})$ is an exact model of Γ, i.e., for each F,

$$\Gamma \vdash F \quad \Leftrightarrow \quad \mathcal{M}(\mathsf{S5}_n), \widetilde{\Gamma} \Vdash F.$$

Indeed, if $\Gamma \vdash F$, then $F \in \widetilde{\Gamma}$ by the definition of $\widetilde{\Gamma}$. By the Truth Lemma in $\mathcal{M}(\mathsf{S5}_n)$, F holds at the world $\widetilde{\Gamma}$. If $\Gamma \not\vdash F$, then, by deductive completeness of Γ, $\Gamma \vdash \neg F$, hence, as before, $\neg F$ holds at $\widetilde{\Gamma}$, i.e., $\mathcal{M}(\mathsf{S5}_n), \widetilde{\Gamma} \not\Vdash F$. $\square$

Theorem 1 shows serious limitations of semantical definitions. Since, intuitively, deductively complete scenarios Γ are exceptions, "most" epistemic situations cannot be defined semantically.

In Section 4.6, we provide a concrete example of an incomplete but meaningful epistemic scenario, a natural variant of the Muddy Children puzzle,

[12]The same criteria hold for any other normal modal logic which has a canonical model.

which, by Theorem 1 does not have a semantical definition, but can nevertheless be easily specified and analyzed syntactically.

In Section 5.1, we consider an example of an extensive game with incomplete epistemic description which cannot be defined semantically, but admits easy syntactical analysis.

4 The Muddy Children puzzle

Consider the 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. Each child sees everybody else's foreheads, but not his 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 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?*

4.1 Standard syntactic formalization

This situation can be described in n-agent epistemic logic $\mathsf{S5}_n$ with modalities $\mathbf{K}_1, \mathbf{K}_2, \ldots, \mathbf{K}_n$ for the children's knowledge and atomic propositions $m_1, m_2, \ldots, m_n$ with m_i stating "child i is muddy."

In addition to general epistemic logic principles $\mathsf{S5}_n$, 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:

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

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

After the father's announcement "some of you are muddy," if a child sees another child not muddy, he knows that he himself is muddy. Since both children announce that they did not know that they were muddy, reflecting on this argument, both children determine that they are muddy and raise their hands in the second round.

This reasoning is quite rigorous and can be formalized in modal epistemic logic. We will present this solution in a more formal setting in Section 4.7.

4.2 Syntactic vs semantic formalizations

The "standard" formalization of Muddy Chidren is by a Kripke model: n-dimensional cube Q_n ([15, 19, 20, 21, 22]). The aforementioned syntactic formalization of the initial configuration of Muddy Children puzzle as MC_n corresponds to Q_n, cf. Section 4.4, and is an example of handling a problem in SEL, via syntactic formalization.

A peculiar feature of both MC_n and Q_n formalizations is that they postulate that no child knows his/her forehead. However, all children eventually know whether they are muddy. This means that in the process if communicating and reasoning, children reverse assumption (2) 'not knowing about themselves.' As we can see in Section 4.7, the syntactic solution of this puzzle does not need assumption (2) and avoids this kind of non-monotonic reasoning.

4.3 Semantic solution

In a semantic solution, the set of assumptions MC_n is replaced by a Kripke model: n-dimensional cube Q_n ([15, 19, 20, 21, 22]). To keep things simple, we consider the case $n = k = 2$. Logical possibilities for the truth value combinations[13] of (m_1, m_2): (0,0), (0,1), (1,0), and (1,1) are declared possible worlds. There are two indistinguishability relations denoted by solid arrows (for agent 1) and dotted arrows (for agent 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. After the father publicly announces $m_1 \vee m_2$, node $(0, 0)$ is no longer possible and model $\mathcal{M}_1$ now becomes common knowledge. Both

[13] 1 standing for 'true' and 0 for 'false'

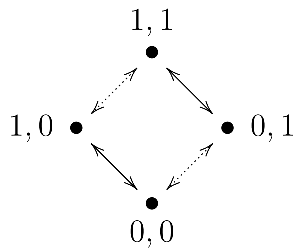

Figure 1: Model Q_2.

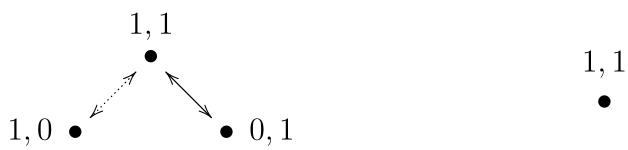

Figure 2: Models $\mathcal{M}_1$ and $\mathcal{M}_2$.

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. 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 $\mathcal{M}_2$. Now both children know that their foreheads are muddy.

4.4 Justifying the model

The semantic solution in Section 4.3 starts with adopting Q_n as an equivalent of a theory MC_n. Can this choice of a 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 $\mathsf{MC}_n(u)$ we understand the Muddy Children scenario with specific distribution of truth values of m_i's corresponding to u:

$$\mathsf{MC}_n(u) = \mathsf{MC}_n \cup \{\pi(u)\}.$$

So, each specific instance of Muddy Children is formalized by an appropriate $\mathsf{MC}_n(u)$.

Theorem 2 *Each instance* $\mathsf{MC}_n(u)$ *of Muddy Children is deductively complete and* (Q_n, u) *is its exact model*

$$\mathsf{MC}_n(u) \vdash F \quad \textit{iff} \quad Q_n, u \Vdash F.$$

Proof.[14] The direction 'only if' claims that (Q_n, u) is a model for $\mathsf{MC}_n(u)$ is straightforward. First, Q_n is an $\mathsf{S5}_n$-model and all principles of $\mathsf{S5}_n$ 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 $\mathsf{MC}_n(u)$ holds at u.

To establish the 'if' direction, we first note that, by Proposition 1, necessitation is admissible in MC_n: for each F,

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

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

for all nodes $u \in Q_n$,

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

and

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

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

The case F is one of the atomic propositions $m_1, m_2, \ldots, m_n$ is trivial since $\mathsf{MC}_n \vdash \pi(u) \rightarrow m_i$, if $u \Vdash m_i$ and $\mathsf{MC}_n \vdash \pi(u) \rightarrow \neg m_i$, 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.

[14] We have chosen to present a syntactic proof of this theorem, partly because it emphasizes the strong and fundamental nature of syntactic formalism. A semantic proof can also be given. It makes use of bounded morphisms, a special case of bi-simulations. This is standard machinery for modal logicians, but is perhaps less familiar to those working in game theory. It is sophisticated and powerful machinery, and should be better known beyond the logic community.

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,

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

By the rules of logic (splitting premises)

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

where $\pi(v)_{-i}$ is $\pi(v)$ without its i-th coordinate[15]. By further reasoning,

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

By necessitation in MC_n, and distributivity,

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

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

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

hence

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

and

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

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,

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

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

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

By contrapositive,

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

By necessitation and distribution,

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

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

By 'knowing about others,' as before,

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

By 'not knowing about himself,' $\mathsf{MC}_n \vdash \neg\mathbf{K}_i m_i$, hence

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

and

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

□

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

given a syntactic description, pick a "natural model"

is intrinsically flawed: by Theorem 1, 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 2 is required.

4.5 Atomic propositions and possible worlds

A possible (and observed) reaction to the criticism that Q_n was adopted as *The Model* of Muddy Children puzzle without a completeness analysis of its syntactic description is

> *It is not much to assume that an agent can figure out that the logical possibilities for the truth values of* m_1, m_2 *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 of truth values of atoms determine truth values of any relevant epistemic sentence.* Without this, we cannot claim that atoms determine a possible world. In the version of the Muddy Children puzzle from Section 4.6, this is not the case.

4.6 Muddy Children lite

Consider a simplified Muddy Children scenario, MClite_2, which has a straightforward syntactic formalization and solution but lacks an exact semantic description. We also show that in MClite_2, possible worlds are not determined by the truth values of atoms.

In Muddy Children lite, condition (2) "not knowing about himself" of the standard Muddy Children puzzle 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 asks: "Do either 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. What do children do?*

The challenge here is to formalize all three stages of this situation: prior to father's announcement, after the announcement, and after the first round of question and answers, and justify children's answers.

The syntactic description of MClite_2 over two-agent $\mathsf{S5}_2$ with atoms m_1 and m_2 is given by

- common knowledge of "knowing about the other"

$$\bigwedge_{i \neq j} [\mathbf{K}_i(m_j) \vee \mathbf{K}_i(\neg m_j)], \quad i, j = 1, 2.$$

- $m_1 \wedge m_2$.

What is the natural epistemic model of MClite_2? Good old Q_2 with world $1,1$ is certainly a model (fewer conditions to check than for MC_2). Here are other models $\mathcal{M}_3$, $\mathcal{M}_4$, and $\mathcal{M}_5$, in which MClite_2 holds at node $1,1$.

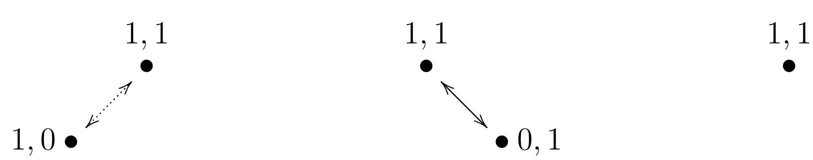

Figure 3: Models $\mathcal{M}_3$, $\mathcal{M}_4$ and $\mathcal{M}_5$.

Incidentally, multiple models appear here not just because some edges in Q_2 are not specified in MClite_2. The problem is deeper: unlike the standard Muddy Children, MClite_2 has models in which truth values of atomic propositions do not determine a possible world. For example, MClite_2 holds at each world of $\mathcal{M}_6$ that has different possible worlds a and b with the same

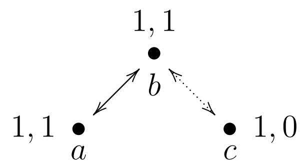

Figure 4: Model $\mathcal{M}_6$.

propositional component $(1,1)$, e.g., $a \Vdash \mathbf{K}_2 m_2$ ("child 2 knows that (s)he is muddy") and $b \Vdash \neg\mathbf{K}_2 m_2$ ("child 2 does not know that (s)he is muddy").

Proposition 3 MClite_2 *is deductively incomplete and hence does not have an exact model.*

Proof. Formula $\mathbf{K}_1 m_1$ holds in $\mathcal{M}_3$ and does not hold in $\mathcal{M}_4$. Therefore, neither $\mathbf{K}_1 m_1$ nor $\neg\mathbf{K}_1 m_1$ is derivable in MClite_2. $\square$

4.7 Syntactic solution of Muddy Children lite

The aforementioned intuitive deductive solution of MC_2 (Section 4.1) fits MClite_2 as well since it does not use condition (2) "not knowing about himself." Moreover, we can avoid knowledge revisions typical for solutions of the usual Muddy Children puzzle which some could find questionable[16].

Here is a syntactic solution of MClite_2. We denote $m_1 \wedge m_2$ as **11**, $\neg m_1 \wedge m_2$ as **01**, $m_1 \wedge \neg m_2$ as **10**, $\neg m_1 \wedge \neg m_2$ as **00**.

The initial description Γ_0 of MClite_2 atop the two-agent epistemic logic $\mathsf{S5}_2$ (cf. Section 2) is a direct formalization of the some default epistemic

[16] In standard MC_2, the ignorance of agents about the state of their foreheads is postulated up front, e.g., $\neg\mathbf{K}_1 m_1$, hence $\mathbf{K}_1\neg\mathbf{K}_1 m_1$, only to be replaced by its negation $\mathbf{K}_1 m_1$ during further exchanges. This constitutes a failure of a natural indefeasibility property of knowledge of $\neg\mathbf{K}_1 m_1$: "there is no further truth which, had the subject known it, would have defeated [subject's] present justification for the belief" [18].

assumptions and the condition "two children have muddy foreheads and each child sees the other child's forehead."[17]

Commonly known assumptions:

1. Knowing About Others (KAO) assumption

$$\mathbf{K}_1 m_2 \vee \mathbf{K}_1(\neg m_2), \quad \mathbf{K}_2 m_1 \vee \mathbf{K}_2(\neg m_1).$$

   This reflects a general idea that by the situation setup, it is common knowledge that each child sees the other child's forehead in the real and all hypothetical situations.

2. The existence of the "real world": the situation occurs under specific truth values of m_1 and m_2. This is reflected in assuming common knowledge that at least one of **11**, **01**, **10**, or **00** should be in Γ_0. This fact is not represented by a formula, but plays a role in determining the updates of agents' knowledge after public announcements.

Specific "real world" condition:

3. **11** ($= m_1 \wedge m_2$) – the true values of m_1, m_2. This corresponds to condition "both children have muddy foreheads."

So,

$$\Gamma_0 = \mathbf{C}(\mathrm{KAO}) \cup \{\mathbf{11}\},$$

with additional meta assumption (2).

When reasoning about logical and epistemic information containing in a given syntactic set Γ, we have to distinguish among the following statements.

- "F holds," meaning that $\Gamma \vdash F$.
- "F is known to agent i," meaning $\Gamma \vdash \mathbf{K}_i F$.

  > Note that "F is not known to agent i" is thus $\Gamma \not\vdash \mathbf{K}_i F$ rather then $\Gamma \vdash \neg\mathbf{K}_i F$. We argue that this is a more appropriate formalization of "F is not known to agent i" in a situation of dynamic changes to the data: it is "i does not yet have enough data to conclude F" ($\Gamma \not\vdash \mathbf{K}_i F$) rather than "$i$

[17] The suggested formalization is one of essentially equivalent ways of specifying this situation syntactically.

knows that he does not and never will know F" ($\Gamma \vdash \neg \mathbf{K}_i F$). This fine distinction allows us to avoid non-monotonicity and knowledge revision during updates.

- "F is common knowledge," meaning $\Gamma \vdash \mathbf{C}F$.[18]

Some immediate observations concerning Γ_0.

- $m_1 \wedge m_2$ (= **11**) holds in Γ_0 but is not known to any agent.

- m_1 is known to agent 2, and m_2 is known to agent 1, but neither is common knowledge.

- both $\neg m_1 \rightarrow \mathbf{K}_2(\neg m_1)$ and $\neg m_2 \rightarrow \mathbf{K}_1(\neg m_2)$ are commonly known. We prove the former, the latter is symmetric. Reason in Γ_0. By KAO, it is common knowledge that $\mathbf{K}_2 m_1 \vee \mathbf{K}_2(\neg m_1)$ and hence that $\neg \mathbf{K}_2 m_1 \rightarrow \mathbf{K}_2(\neg m_1)$. By reflection, it is common knowledge that $\neg m_1 \rightarrow \neg \mathbf{K}_2 m_1$. By logic reasoning, it is common knowledge that $\neg m_1 \rightarrow \mathbf{K}_2(\neg m_1)$.

The public announcements, questions, and answers lead to two consecutive updates of Γ_0 to Γ_1, and then to Γ_2. Unlike the traditional Muddy Children puzzle, we do not revise knowledge here: both updates are monotonic, agents just acquire new information without rejecting their existing data.

Update 1. Public announcement "*some of the children are muddy*" amounts to making "not **00**" common knowledge.

$$\Gamma_1 = \Gamma_0 + \mathbf{C}(\neg \mathbf{00}).$$

This rules out the option "both are not muddy" ($\mathbf{00} \in \Gamma_1$) for both children.

Update 2. Father's question "Do any of you know that you have mud on your forehead?" and the negative replies reveal that neither of the children can conclude that he is muddy, i.e.,

$$\Gamma_1 \not\vdash \mathbf{K}_1 m_1 \quad \text{and} \quad \Gamma_1 \not\vdash \mathbf{K}_2 m_2.$$

This information is incompatible with situations $\mathbf{01} \in \Gamma_1$ and $\mathbf{10} \in \Gamma_1$. Indeed, suppose **01** ($= \neg m_1 \wedge m_2$) were in Γ_1 instead of **11**. Since Γ_1 contains

$$\neg m_1 \rightarrow \mathbf{K}_2(\neg m_1),$$

[18] This means Γ proves each of F, $\mathbf{E}F$, $\mathbf{E}^2 F$, $\mathbf{E}^3 F$, ….

we would also have $\mathbf{K}_2(\neg m_1)$. By the first update, $\mathbf{K}_2\neg(\neg m_1 \wedge \neg m_2)$ and, by the rules of logic, $\mathbf{K}_2(\neg m_1 \rightarrow m_2)$, hence

$$\mathbf{K}_2(\neg m_1) \rightarrow \mathbf{K}_2 m_2.$$

From this we would derive

$$\mathbf{K}_2 m_2,$$

i.e., child 2 would know that he was muddy, which is inconsistent with his/her answer. Case $\mathbf{10} \in \Gamma_1$ is symmetric.

This makes it common knowledge that neither **01** nor **10** is in Γ_1 and hence $\mathbf{11} \in \Gamma_1$ is the only remaining option[19]. This leads to the second update

$$\Gamma_2 = \Gamma_1 + \mathbf{C}(\mathbf{11}),$$

i.e., it becomes common knowledge that both children are muddy.

So the syntactic solution of MClite_2 is given by the sequence of syntactic sets Γ_0, Γ_1, and Γ_2 that provide (partial) descriptions of possible worlds at the corresponding moments. These descriptions themselves contain all relevant information and ignore irrelevant, higher-order epistemic assertions.

4.8 Discussion

The idea of this syntactic solution does not differ much from the aforementioned semantic solution of MC_2: both are based on consecutive elimination of logical possibilities by some kind of epistemic reasoning and meta-reasoning. However, the format of the Kripke/partition model forces operating with individual possible worlds (nodes of the model) which are overspecified.

> *For example, in MClite_2, one logical possibility $m_1 \wedge m_2$ corresponds to infinitely many (continuum) nonconstructive maximal consistent sets each of which can be a possible world. This renders semantic analysis at the level of possible worlds unfeasible.*

So, if we want to go beyond complete epistemic scenarios, we need a mathematical apparatus to handle classes of models, not just single models. The format of syntactic specifications is a viable candidate for such an apparatus.

[19] Technically, this argument uses meta assumption (2).

Note that no belief revision occurs in the syntactic solution of MClite_2: all updates are consistent with the original data.

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, 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 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.

5 Syntactic Epistemic Logic and games

Aumann's definition of an extensive form game ([7]) is based on the notion of a partition structure with a condition that the model should itself be common knowledge. In this paper, we focus on the non-probabilistic case.

We argue that the single-model semantic definition of a game is too restrictive and needs to be extended.

Let us recap basic terminology ([7]). An extensive game consists of the following components.

- A finite set $N = \{1, 2, \ldots, n\}$ of players.
- A finite rooted tree.
- A player function P that assigns a player (who makes a move) to each nonterminal node.
- For each player i, a payoff function u_i defined on terminal nodes.

The root node is the starting point of the game. At any nonterminal node v, player $P(v)$ chooses one of the successor nodes (a move). A strategy s_i of player i is a choice of moves at each node assigned to player i. A strategy profile $s = (s_1, \ldots, s_n)$ is an n-tuple of strategies of players $1, 2, \ldots, n$.

An extensive-form game is a game tree and a partition structure $\mathcal{A}$ (cf. Section 3.1) which are commonly known (and a "real state" which makes the game concrete). It is presumed that $\mathcal{A}$ codifies possible worlds 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.

Consider Game 1 with the game tree as 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 ratio-

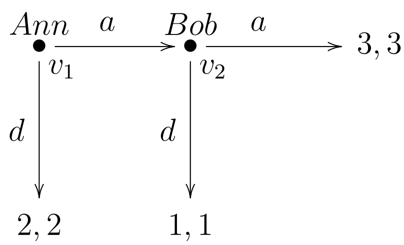

Figure 5: Game 1.

nal; she anticipates Bob's move and herself moves across, hence the "backward induction" solution – strategy profile (a, a) which is common knowledge.

The logic language of this game consists of propositional variables p_A and p_B symbolizing Ann's and Bob's moves across, logical connectives, and knowledge modalities $\mathbf{K}_A$, $\mathbf{K}_B$. The syntactic formalization over $\mathsf{S5}_2$ can be reduced to the following set of formulas:

1. $\mathbf{C}p_B$ - "it is common knowledge that Bob is rational hence plays across";

2. $\mathbf{C}(\mathbf{K}_A p_B \rightarrow p_A)$ - "it is common knowledge that since Ann is rational, if Ann knows that Bob plays across, she plays across as well";

Let Δ_1 be the set of assumptions 1 and 2. Since $\mathbf{C}p_B$ implies $\mathbf{C}\mathbf{K}_A p_B$, by easy modal reasoning we conclude $\mathbf{C}p_A$. This yields that

$$\Delta_1 = \mathbf{C}(p_B \wedge p_A).$$ [20]

The corresponding Kripke model for Δ_1 is $\mathcal{M}_7 = \langle W, R_A, R_B, \Vdash \rangle$ that has only one node (a, a), R_A and R_B are reflexive, and p_A and p_B hold.

$$
\begin{array}{c} p_A, p_B \\ \bullet \\ (a, a) \end{array}
$$

Figure 6: Model $\mathcal{M}_7$.

Δ_1 is similar to the set from Example 2.3 and $\mathcal{M}_7$ is its adequate model representation: for each formula F

$$\Delta_1 \vdash F \quad \text{iff} \quad \mathcal{M}_7 \Vdash F.$$

5.1 Partial knowledge is not definable semantically

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: *mutual knowledge of rationality*, i.e., both Ann and Bob know that both players are rational. Note that Game 2 does not specify second-order knowledge of rationality, e.g., whether Bob knows that Ann knows that Bob is rational, etc. The outcome of Game 2 is the same as that of Game 1: Ann and Bob play across; this can be established by an easy logical reasoning.

Let us consider the syntactic formalization of Game 2.

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

$$\mathbf{K}_A p_B.$$

2. *Ann is rational* is formalized as *if Ann knows that Bob plays across, she plays across* $\mathbf{K}_A p_B \rightarrow p_A$. *Bob knows that Ann is rational* is then

$$\mathbf{K}_B(\mathbf{K}_A p_B \rightarrow p_A).$$

3. Standard 'measurability' condition *knowing their own moves*:

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

The formalization of Game 2 is the set Δ_2 consisting of principles 1 – 3.

Proposition 4 *Game 2 (formalized as Δ_2) is not deductively complete and hence does not have a single-model semantic description.*

[20]Indeed, condition $\mathbf{C}(\mathbf{K}_A p_B \rightarrow p_A)$ immediately follows from $\mathbf{C} p_A$.

Proof. We claim that $\mathbf{K}_B p_A$ (i.e. 'Bob knows that Ann plays across') is independent of Δ_2. Indeed, $\Delta_2 \nvdash \neg\mathbf{K}_B p_A$ since both Δ_2 and $\mathbf{K}_B p_A$ hold in $\mathcal{M}_7$. To show that $\Delta_2 \nvdash \mathbf{K}_B p_A$, consider model $\mathcal{M}_8$. The solid arrow stays

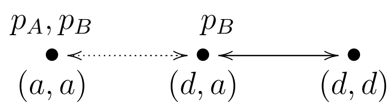

Figure 7: Model $\mathcal{M}_8$.

for R_A and the dotted arrow represents R_B. It is easy to check that Δ_2 holds at (a, a). Let us check (2). Formula $\mathbf{K}_A p_B$ does not hold at (d, a) and (d, d) since p_B does not hold at (d, d) and these two nodes are R_A-equivalent. Hence $\mathbf{K}_A p_B \to p_A$ holds everywhere in the model, hence $\mathbf{K}_B(\mathbf{K}_A p_B \to p_A)$ holds at (a, a).

Furthermore, $\mathbf{K}_B p_A$ does not hold at (a, a) since p_A does not hold at (d, a) and these two nodes are R_B-equivalent. Therefore, $\Delta_2 \nvdash \mathbf{K}_B p_A$.

$\square$

Model $\mathcal{M}_8$ can be represented as a partition structure $\mathcal{A}_8$ that has three possible worlds,

$$\Omega = \{\omega_1, \omega_2, \omega_3\}$$

where ω_1 corresponds to profile (a, a), ω_2 to profile (d, a), and ω_3 to profile (d, d); p_A holds at ω_1, p_B holds at ω_1 and ω_2. Partitions for Ann and Bob are

$$\mathcal{K}_A = \{\{\omega_1\}, \{\omega_2, \omega_3\}\}, \qquad \mathcal{K}_B = \{\{\omega_1, \omega_2\}, \{\omega_3\}\},$$

and the real state is ω_1. In $\mathcal{K}_A$, the partition cell containing ω_1 is a singleton at which p_A holds, 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, therefore, $\neg\mathbf{K}_B p_A$ holds at ω_1, i.e., Bob does not know Ann's move.

5.2 Asymmetric knowledge is not definable either

Consider Game 3 that features the same players Ann and Bob, and the same game tree as Game 1 and Game 2, but with a different epistemic condition: *Ann is rational and knows that Bob is rational.* Note that Game 3 does not specify whether Bob knows that Ann is rational, not to mention higher-order epistemic assertions of the type 'Ann knows that Bob does not know that she is rational,' etc. The outcome of Game 3 is the same as that of Game 1 and Game 2, namely, Ann and Bob play across.

Here is a syntactic formalization of Game 3.

1. *Ann knows that Bob is rational* is formalized as

$$\mathbf{K}_A(p_B).$$

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

$$\mathbf{K}_A(p_B) \rightarrow p_A.$$

3. Standard 'measurability' condition *knowing their own moves*:

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

The formalization of Game 3 is the set Δ_3 consisting of principles 1 – 3. Obviously, this is equivalent to the set Δ'_3: [21]

$$\{\mathbf{K}_A(p_A),\ \mathbf{K}_A(p_B),\ \mathbf{K}_B(p_B)\}.$$

It is easy to see that $\mathbf{K}_B(p_A)$ is independent of Δ'_3: in model $\mathcal{M}_7$, Δ'_3 and $\mathbf{K}_B(p_A)$ are true, model $\mathcal{M}_9$ at (a,a) makes Δ'_3 true and $\mathbf{K}_B(p_A)$ false.

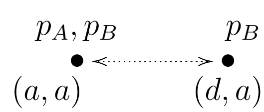

Figure 8: Model $\mathcal{M}_9$

Model $\mathcal{M}_9$ can be represented as a knowledge structure $\mathcal{A}_9$ 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), p_A holds at ω_1, and p_B holds at both states. 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, $\neg\mathbf{K}_B(p_A)$ holds in $\mathcal{A}_9$, but does not follow from Game 3. So $\mathcal{A}_9$ is not an adequate semantic formalization of Game 3 either.

[21] All principles of Δ'_3 logically follow from 1–3 and all 1–3 can be derived from Δ'_3.

Proposition 5 *Game 3 (formalized by Δ_3) is not deductively complete and hence does not have an exact model.*

6 Incomplete Scenarios

How typical are deductively incomplete epistemic scenarios? We argue informally that this is the rule rather than the exception: it appears that unless common knowledge of the basic assumptions is postulated, some higher-order epistemic assertions remain independent. Epistemic conditions more flexible than CKGR (mutual knowledge of rationality, asymmetric epistemic assumptions when one player knows more than the other, etc.) lead to semantical undefinability.

Of course, common knowledge of epistemic conditions does not alone yield deductive completeness. The Muddy Children lite scenario has commonly known epistemic assumptions but is, nevertheless, incomplete.

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

7 Complete Scenarios

By Theorem 1, deductively incomplete scenarios do not have single-model semantic characterizations. The principal 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 1 are instances of the canonical model $\mathcal{M}(\mathsf{S5}_n)$ at the node $\widetilde{\Gamma}$ corresponding to Γ. This generic solution is, however, not satisfactory because of the highly nonconstructive nature of the canonical model $\mathcal{M}(\mathsf{S5}_n)$.

As was shown in [8], the canonical model $\mathcal{M}(\mathsf{S5}_n)$ for any $n \geq 1$ has continuum-many possible worlds even with just one propositional letter. This alone renders models $(\mathcal{M}(\mathsf{S5}_n), \widetilde{\Gamma})$ not knowable under any reasonable meaning of "known." The canonical model for $\mathsf{S5}_n$ is just too large to be considered known and hence does not *a priori* satisfy the knowability of the model requirement.

So, we have to admit that the question about existence of an epistemologically acceptable ("known") model for a given deductively complete set Γ remains open.

A lazy cardinality analysis shows that "most" deductively complete Γ's don't have knowable single-model characterizations: there are continuum-many such Γ's, basically each possible world in $\mathcal{M}(\mathsf{S5}_n)$. However, there are only countably-many structures/models that can be specified by efficient descriptions in a countable language (finite, computable, computably-enumerable, arithmetical, etc.). This argument suggests that the question of semantical definability of deductively complete scenarios, perhaps, requires a case-by-case consideration.

7.2 Specifications via finite models and SEL

If an epistemic scenario can be naturally formalized by a manageable model, we have a win-win situation. Each model $(\mathcal{M}, u)$ specifies a sets of formulas

$$\Gamma(u) = \{F \mid \mathcal{M}, u \Vdash F\}.$$

For a finite structure $\mathcal{M}$[22], $\Gamma(u)$ is deductively closed and linear-time decidable. Indeed, we have to check that the validity at u is linear-time decidable. The validity of atoms is given, and with each step of formula-building we have to check the bit values of one or two subformulas at the fixed finite set of nodes which amounts to a fixed finite amount of checkups.

As we can see, the set of formulas $\Gamma(u)$ provides an efficient syntactic description of the epistemic situation specified by model $(\mathcal{M}, u)$. So, semantic specifications by finite models are acceptable in SEL and even desirable.

7.3 Complexity considerations

One could argue that in a complete scenario, it is natural to reason in terms of models rather than in terms of logical conditions. We believe this is not necessarily the case.

[22]A relational structure is finite if both the set of states and the evaluation of variables are finite (e.g., the evaluation $u \Vdash p_i$ holds only for a finite set of variables p_i).

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

Consider a real-life situation, say the game of poker. One can show that its natural syntactic description in epistemic logic is deductively complete and hence admits a model characterization. Moreover, it has a natural finite model of the type given in Chapter 2.1 of [15], 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. So, the number of nodes in the model exceeds the number of stars in the observable universe. This yields that explicit formalization of the model not practical. Players use a concise syntactic description of poker in the natural language, which can also be formalized in epistemic logic[23] in a feasible way.

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

8 Some Practical Observations

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. Common knowledge of the game and rationality assumption CKGR often yield deductive completeness. For example, in perfect information games, CKGR leads to deductive completeness and non-problematic singleton models. However, CKGR is considered too restrictive: players might not have complete and equal information about the game and each other, there might be a certain level of ignorance and/or secrecy, etc. In these cases, CKGR does not hold, which leaves the door open for semantically non-definable scenarios.

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 a "natural model" $\mathcal{M}$ of G, normally overspecified, we try not to use features of the model that are not supported

[23] perhaps first-order epistemic logic

by G. If we conclude a property P by such self-restricted reasoning about the model, then P indeed logically follows from G.

> *This situation resembles Geometry, in which we reason about triangles, circles, etc., but in the background have a rigorous system of postulates. We are trained not to venture beyond these postulates even in informal reasoning.*

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

9 Syntactic Epistemic Logic suggestions

The Syntactic Epistemic Logic suggestion, in brief, is *to make the syntactic formalization of an epistemic scenario its formal specification*.

Since SEL accommodates constructive model reasoning (Section 7.2), we speak about extending the scope of scientific epistemology. SEL offers a remedy for two principal weaknesses of the traditional semantic approach.

1. SEL suggests a way to handle incomplete scenarios which, however, have reasonable syntactic descriptions (cf. MClite_2, Game 2, Game 3, etc.).

2. SEL provides a scientific framework for resolving the tension, identified by R. Aumann [9], between a syntactic description and its hand-picked model: formalize the former and establish completeness.

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

Within Syntactic Epistemic Logic, we offer a new definition of an extensive form game,

> *a game tree plus a syntactic description of epistemic conditions*,

which is more general than the knowledge model definition. Numerous games with deductively incomplete descriptions can now be formalized (cf. Game 2, Game 3) and studied.

This can help to extend Epistemic Game Theory from CKGR towards capturing more general 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 model characterizations. However, if such a scenario allows an adequate syntactic formulation, it can be handled by a variety of mathematical tools, including reasoning about its models.

If an epistemic scenario is given syntactically but represented by an epistemic model, it is prudent to establish their equivalence.

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. The case study of syntactic updates is presented in the syntactic solution of the Muddy Children lite puzzle in section 4.7.

10 Acknowledgements

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

References

[1] I. Arieli and R. Aumann. *The logic of backward induction.* Available at SSRN 2133302, 2012.

[2] S. Artemov. Explicit provability and constructive semantics. *Bulletin of Symbolic Logic*, 7(1):1–36, 2001.

[3] S. Artemov. Justified common knowledge. *Theoretical Computer Science*, 357(1-3):4–22, 2006.

[4] S. Artemov. 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, Springer, pp. 487–507, 2014.

[5] S. Artemov and M. Fitting. Justification Logic. *Stanford Encyclopedia of Philosophy*, 2011.

[6] R. Aumann. Agreeing to disagree. *The Annals of Statistics*, 4(6):1236–1239, 1976.

[7] R. Aumann. Backward induction and common knowledge of rationality. *Games and Economic Behavior*, 8(1):6–19, 1995.

[8] R. Aumann. Interactive epistemology I: Knowledge. *International Journal of Game Theory*, 28:263–300, 1999.

[9] R. Aumann. Epistemic Logic: 5 Questions. Vincent F. Hendricks and Olivier Roy, eds. Automatic Press/VIP, pp. 21–33, 2010.

[10] R. Aumann and A. Brandenburger. Epistemic conditions for Nash equilibrium. *Econometrica*, 63:1161–1180, 1995.

[11] P. Blackburn, M. de Rijke and Y. Venema. *Modal Logic.* Volume 53 of Cambridge Tracts in Theoretical Computer Science, 2001.

[12] P. Blackburn and J. van Benthem. Modal logic: A semantic perspective. *Handbook of Modal Logic.* pp.1–84. Studies in Logic and Practical Reasoning 3, Elsevier, 2007

[13] A. Brandenburger. *The Language of Game Theory: Putting Epistemics Into the Mathematics of Games.* World Scientific Publishing Company, 2014.

[14] A. Chagrov and M. Zakharyaschev. *Modal Logic.* Volume 35 of Oxford Logic Guides, 1997

[15] R. Fagin, J. Halpern, Y. Moses and M. Vardi. *Reasoning About Knowledge.* MIT Press, 1995.

[16] M. Fitting. Modal proof theory. *Handbook of Modal Logic.* pp.85–138. Studies in Logic and Practical Reasoning 3, Elsevier, 2007

[17] J. Hintikka. *Knowledge and Belief.* Ithaca, N.Y.,Cornell University Press, 1962.

[18] K. Lehrer and T. Paxson. Knowledge: undefeated justified true belief. *The Journal of Philosophy*, 66:1–22, 1969.

[19] J.J.C. Meyer and W. van der Hoek. *Epistemic Logic for AI and Computer Science.* Volume 41 of Cambridge Tracts in Theoretical Computer Science, 1995.

[20] M. Osborne and A. Rubinstein. *A Course in Game Theory.* MIT Press, 1994.

[21] M. Pauly, and W. van der Hoek. Modal logic of information and games. *Handbook of Modal Logic.* pp.1077–1148. Studies in Logic and Practical Reasoning 3, Elsevier, 2007

[22] J. van Benthem. *Logic in Games.* MIT Press, 2014.