Subjective situations

Antonio Moreno[1], Ulises Cortés[2], Ton Sales[2]

[1] Departament d'Enginyeria Informàtica - Universitat Rovira i Virgili (URV)
Carretera de Salou, s/n. 43006-Tarragona, Spain - amoreno@etse.urv.es
[2] Dep. de Llenguatges i Sistemes Informàtics - Universitat Politècnica de Catalunya
C/Jordi Girona, 1-3. 08034-Barcelona, Spain - {ia,sales}@lsi.upc.es

Abstract. The beliefs of the agents in a multi-agent system have been formally modelled in the last decades using *doxastic logics*. The *possible worlds model* and its associated *Kripke semantics* provide an intuitive semantics for these logics, but they commit us to model agents that are *logically omniscient*. We propose a way of avoiding this problem, using a new kind of entities called *subjective situations*. We define a new doxastic logic based on these entities and we show how the belief operators have some desirable properties, while avoiding logical omniscience. A comparison with two well-known proposals (Levesque's *logic of explicit and implicit beliefs* and Thijsse's *hybrid sieve systems*) is also provided.

1 Introduction

In the last decade *doxastic modal logics* have been considered the most appropriate formal tool for modelling the beliefs of the agents composing a multi-agent system ([1]). The standard way of providing a meaning to the modal formulas of these logics is to use the *possible worlds model* ([2]) and its associated *Kripke semantics* ([3]). This semantics is quite natural and intuitive, but it is well known that the agents modelled in this framework are *logically omniscient* ([4]), because they must believe all valid formulas and all the logical consequences of their beliefs. Therefore, this semantics is unsuitable to model the beliefs of realistic, non-ideal agents. The aim of our work is to provide a plausible way of modelling the beliefs of non-logically omniscient agents, while keeping the essence and the beauty of the possible worlds model and the Kripke semantics.

This article[3] is structured as follows. In section 2 we give an intuitive explanation of our approach to the logical omniscience problem, which is based in a new kind of entities called *subjective situations*. In a nutshell, a subjective situation is the perception that an agent has of a certain state of affairs. These situations, as will be explained below, will take the role of possible worlds. In section 3, a formalization of subjective situations in the framework of doxastic propositional logic is made. Section 4 is devoted to a study of the behaviour of the modal belief operators. It is shown how their properties do indeed correspond with our intuitions about what should be an adequate formalization of

[3] This research has been supported by the CICYT project *SMASH: Multi-agent systems and its application to hospital services (TIC96-1038-C04-04).*

the doxastic attitude of a non-ideal, non-logically omniscient agent. In section 5, a comparison of our proposal with two well-known approaches (Levesque's *logic of explicit and implicit beliefs* ([5]) and Thijsse's *hybrid sieve systems* ([6])) is performed. The paper finishes with a brief summary and the bibliographical references.

2 Considering Subjective Situations

One of the ways of solving the logical omniscience problem is to revise the concept of possible world[4]. In the classical view, a possible world is a logical model (*e.g.* in a propositional setting, a possible world is an interpretation of the basic propositions). Some authors have tried to alleviate logical omniscience by changing this notion; for instance, a possible world may be seen as a situation in which some basic propositions may be considered true and some basic propositions (not necessarily different from the previous ones) may be considered false ([5]), or it may be represented as a *partial* assignment of truth values to the basic propositions ([6]) or simply as the set of formulas that are assumed to be true in it ([8], [9]).

Regardless of the way the concept of possible world is modified, there is something that does not change: the formal representation of a possible world is not related in any way with the notion of *agent*. Thus, it may be said that all the approaches in the literature present an *objective* view of what a possible world is (*i.e.* a world is the same for all the agents). In a standard *Kripke structure*, the only item that depends on each agent is the *accessibility relation* between possible worlds. The traditional meaning attached to the accessibility relation R_i of an agent i is that it represents the uncertainty that agent i has about the situation in which it is located (*e.g.* $(w_0 R_5 w_1)$ means that agent 5 cannot distinguish between worlds w_0 and w_1, see [1]). This situation is quite peculiar, because the formulas that are true in two worlds that are linked by an accessibility relation are, in principle, totally unrelated.

Our proposal may be motivated by the following scenario. Imagine two people (α and β) that are watching a football match together. In a certain play, the referee awards a penalty kick. α thinks that the referee is right, because a fault was made inside the penalty area (let's represent this fact with proposition P); at the same time, β is thinking that the referee was wrong, because, in its perception of the situation, the fault was made just an inch outside the penalty area. How can this situation (and the beliefs of the two agents) be formally represented? This state of affairs may not be described with an assignment of truth values to the basic propositions. The situation (s) is obviously the same for the two agents α and β (they are watching the same match together). From α's point of view, the description of s should make true proposition P; however, from β's

[4] Other alternatives, such as getting rid of the Kripke semantics or changing the underlying consequence relation, have also been considered in the literature. A detailed review of the most interesting proposals that have been made to solve the logical omniscience problem may be found in [7].

perspective, in the present situation P should be considered false. Obviously, there would be many aspects of s in which α and β would agree; *e.g.* both of them would consider that the proposition representing the fact *"We are watching a football match on TV"* is true in s.

As far as beliefs are concerned, we argue that, in this situation, α should be capable of stating that $B_\alpha P$ (α has seen the fault and has noticed that it was made inside the penalty area). It would not seem very acceptable a situation in which α perceived the fault to have been made inside the penalty area and defended that it didn't believe that a penalty kick should have been awarded. It also seems reasonable to say that α cannot fail to notice that it believes that the fault was made inside the penalty area; thus, α may also assert in s that $B_\alpha B_\alpha P$. In a similar way, in this situation β cannot state that $B_\beta P$ (β cannot defend that it believes that the referee is right, in a situation in which it perceived the fault to have been made outside the penalty area).

In our framework we want to include the intuition that agents are smart enough to know that other agents may not perceive reality in the same way as they do. In the previous example, without further information (*e.g.* β shouting *Penalty!*), α should not be capable of supporting (or rejecting) that $B_\beta P$; analogously, β could not affirm (or deny) that $B_\alpha P$.

A final reflection on the meaning of the accessibility relation between situations for agent i (R_i) is necessary. It will be assumed that an agent cannot have any doubts about its own perceptions and beliefs in a given state. *E.g.* if, in situation s, α looks at the match and thinks P, then it surely must realise this fact and believe P in s (and even believe that it believes P, were it to think about that). Thus, if R_α links s with all those situations that α cannot tell apart from s, it must be the case that α also perceives P as true in all those states as well (otherwise, those states would be clearly distinguishable by α, because in some of them it would support P whereas in some of them P would be rejected). The only uncertainty that α may have is *about the perception of s by the other agents*. In the example, α does not know whether it is in a situation in which β supports P or in a situation in which β rejects P. Therefore, α's accessibility relation must reflect this uncertainty.

Summarising, the main points that have been illustrated with the previous discussion are the following:

- A situation may be considered not as an entity that may be objectively described, but as a piece of reality that may be perceived in different ways by different agents.
  Thus, it is necessary to think of a *subjective* way of representing each situation, in which each agent's point of view is taken into account. In the previous example, the description of s should include the fact that α is willing to support P, whereas β isn't.
- The beliefs of an agent in each situation also depend on each agent's point of view.
  In the situation of the example, $B_\alpha P$ would hold from α's perspective, whereas it would not be either supported or rejected by β. Thus, we ar-

gue that it does not make sense to ask whether $B_\alpha P$ holds in s or not; that question must be referred to a particular agent's point of view.

– The interpretation of the meaning of each agent's accessibility relation is slightly different from the usual one.
Each accessibility relation R_i will keep its traditional meaning, *i.e.* it will represent the uncertainty of agent i with respect to the situation in which it is located. However, our intuition is that an agent may only be uncertain about the other agents' perception of the present state, not about its own perception.

3 Formalizing Subjective Situations

These intuitive ideas are formalized in the *structures of subjective situations*:

Definition 1 (Structure of Subjective Situations)

An structure of subjective situations E is a tuple

$$< S, R_1, ..., R_n, \mathcal{T}_1, ...\mathcal{T}_n, \mathcal{F}_1, ..., \mathcal{F}_n >, \textit{ where}$$

– *S is the set of possible situations.*
– *R_i is the accessibility relation between situations for agent i.*
– *$\mathcal{T}_i$ is a function that returns, for each situation s, the set of propositional formulas that are perceived as true by agent i in s.*
– *$\mathcal{F}_i$ is a function that returns, for each situation s, the set of propositional formulas that are perceived as false by agent i in s.*

$\mathcal{E}$ is the set of all structures of subjective situations.

The presence of $\mathcal{T}_i$ and $\mathcal{F}_i$ allows agent i to consider *partial* situations (those in which i does not have any reason to support or to reject a given formula) as well as *inconsistent* situations (those in which i may have reasons to support and to reject a given formula).

The accessibility relation between situations for agent i has to reflect its uncertainty about the way that the actual situation is perceived by the other agents. Thus, R_i has to link all those states that i perceives in the same way but that may perceived in different ways by other agents. This intuition is formalized in the following condition:

Definition 2 (Accessibility Relations)

$$\forall s, t \epsilon S, (sR_it) \textit{ if and only if } (\mathcal{T}_i(s) = \mathcal{T}_i(t)) \textit{ and } (\mathcal{F}_i(s) = \mathcal{F}_i(t))$$

This condition forces each R_i to be an equivalence relation, thereby linking this approach with the classical $S5$ modal system, in which this condition also holds. In $S5$ the presence of this condition makes true axiom 4 (positive introspection), axiom 5 (negative introspection) and axiom T (the axiom of knowledge); the modal operators of the system proposed in this article will have similar properties, as will be shown in section 4.

3.1 Satisfiability Relations

A simplified version of the doxastic propositional language for n agents is considered, as shown in the following definition:

Definition 3 (Doxastic Modal Language $\mathcal{L}$)
Consider a set of modal belief operators for n agents (B_1, ..., B_n). $\mathcal{L}$ is the language formed by all propositional formulas (built in the standard way from a set $\mathcal{P}$ of basic propositions and the logical operators $\neg, \vee, \wedge, \rightarrow$), preceded by a (possibly empty) sequence of (possibly negated) modal operators. $\mathcal{L}_{PC}$ is the subset of $\mathcal{L}$ that contains those formulas that do not have any modal operator. The modal formulas of $\mathcal{L}$ are called linearly nested.

Thus, $\mathcal{L}$ contains formulas such as P, B_3Q, $B_1B_5(R \vee T)$, $B_3\neg B_2S$ and $\neg B_1B_1\neg T$, but it is not expressive enough to represent formulas such as $(B_2P \rightarrow B_3Q)$ or $(P \vee B_5Q)$. In most practical applications, an agent in a multi-agent system will only need to represent what it believes (or not) to be the case in the world and what it believes (or not) that the other agents believe (or not). This is just the level of complexity offered by linearly nested formulas.

Following the previously outlined intuitions, two relations (of satisfiability, $\models_i$, and unsatisfiability, $\dashv\!|_i$) between situations and formulas will be defined *for each agent* i. Given an structure of subjective situations E and a situation s, $E, s \models_i \phi$ should hold whenever agent i has some reason to think that ϕ is true in situation s. Similarly, $E, s \dashv\!|_i \phi$ should hold whenever agent i has some reason to reject ϕ in situation s.

Notice that $E, s \not\models_i \phi$ should not imply that $E, s \dashv\!|_i \phi$ (i.e. i not having any reason to support ϕ does not mean that it must have reasons to reject it). In the same spirit, $E, s \models_i \phi$ should not imply $E, s \not\dashv\!|_i \phi$ (agent i could have reasons both to support and to reject a certain formula in a given situation).

The clauses that define the behaviour of these relations are the following:

Definition 4 (Relations $\models_i$ and $\dashv\!|_i$)

– $\forall E \epsilon \mathcal{E}, \forall s \epsilon S, \forall \mathit{agent}\, i, \forall \phi \epsilon \mathcal{L}_{PC}$

$$E, s \models_i \phi \Leftrightarrow \phi \epsilon \mathcal{T}_i(s)$$
$$E, s \dashv\!|_i \phi \Leftrightarrow \phi \epsilon \mathcal{F}_i(s)$$

– $\forall E \epsilon \mathcal{E}, \forall s \epsilon S, \forall \mathit{agents}\, i, j,\ \forall \phi \epsilon \mathcal{L}$

$$E, s \models_i B_j\phi \Leftrightarrow \forall t \epsilon S\ ((sR_it)\ \mathit{implies}\ E, t \models_j \phi)$$
$$E, s \dashv\!|_i B_j\phi \Leftrightarrow \exists t \epsilon S\ ((sR_it)\ \mathit{and}\ E, t \dashv\!|_j \phi)$$

– $\forall E \epsilon \mathcal{E}, \forall s \epsilon S, \forall \mathit{agents}\, i, j,\ \forall \phi \epsilon \mathcal{L}$

$$E, s \models_i \neg B_j\phi \Leftrightarrow E, s \dashv\!|_i B_j\phi$$
$$E, s \dashv\!|_i \neg B_j\phi \Leftrightarrow E, s \models_i B_j\phi$$

Thus, in a given situation s, agent i supports that j believes ϕ just in case j supports ϕ in all the situations that are considered possible by i in s. Similarly, i may reject the fact that j believes ϕ if it may think of a possible situation in which j rejects ϕ.

4 Properties of the Belief Operators

The definition of a structure of subjective situations, the fact that the accessibility relations are equivalence relations and the clauses that describe the behaviour of the satisfiability (and unsatisfiability) relations imply that the modal belief operator of each agent i has several interesting logical properties (that, in our opinion, make it an appropriate operator to model the notion of belief for a non-ideal agent). Some of these properties are described in this section[5].

General Results

Proposition 1 (Lack of Logical Omniscience)

In the framework of subjective situations, none of the following forms of logical omniscience ([1], [7]) holds:

- *Full logical omniscience.*
- *Belief of valid formulas.*
- *Closure under logical implication.*
- *Closure under logical equivalence.*
- *Closure under material implication.*
- *Closure under conjunction.*
- *Weakening of beliefs.*
- *Triviality of inconsistent beliefs.*

These properties do not hold because $\mathcal{T}_i$ and $\mathcal{F}_i$ are defined on sets of (arbitrary) formulas (not on basic propositions). It is possible to impose any of the above properties on the belief operators by requiring these sets of formulas to satisfy some conditions (e.g. if $(\phi \wedge \psi) \epsilon \mathcal{T}_i(s)$ implies that $\phi \epsilon \mathcal{T}_i(s)$ and $\psi \epsilon \mathcal{T}_i(s)$, then the set of beliefs of agent i would be closed under conjunction).

Proposition 2 (Relation between $\models_i$ and $=\!|_i$)

The following properties do not hold:

$$E, s \not\models_i \phi \textit{ implies } E, s =\!|_i \phi$$
$$E, s \models_i \phi \textit{ implies } E, s \not=\!|_i \phi$$

This result is due, as suggested above, to the fact that an agent may have reasons to support *and/or* deny a formula in a given situation.

[5] The proofs of the following propositions are quite straightforward; they have not been included in the final version of the paper due to lack of space.

Results on Positive Introspection

Proposition 3 (Single-agent Positive Introspection)

$$E, s \models_i B_i\phi \textit{ implies } E, s \models_i B_iB_i\phi$$

This proposition states that axiom 4 (the classical axiom of positive introspection) holds for each belief operator B_i (*i.e.* every agent has introspective capabilities on its own positive beliefs).

Proposition 4 (Inter-agent Positive Introspection)

$$E, s \models_i B_j\phi \textit{ implies } E, s \models_i B_jB_j\phi$$

This is a more general result (the previous proposition reflects the case $i = j$). It states that each agent is aware of the fact that the other agents also have introspective capabilities.

Proposition 5 (Multi-agent Positive Introspection)

It does not hold (for three different agents i, j, k) that

$$E, s \models_i B_j\phi \textit{ implies } E, s \models_i B_kB_j\phi$$

This proposition states a negative result. It is telling that even if agent i has reasons to support that j believes something, that is not enough for i to think that any other agent will have that belief. This proposition is essentially expressing the uncertainty of agent i about the beliefs of a different agent k.

Proposition 6 (Generation of Positive Beliefs)

$$E, s \models_i B_j\phi \textit{ implies } E, s \models_i B_iB_j\phi$$

If an agent has reasons to support a certain belief, then that belief will be included in its belief set. This proposition is similar to the following one, that characterizes an agent's set of propositional positive beliefs in a given state.

Proposition 7 (Characterization of Positive Beliefs)

For any propositional formula ϕ,

$$E, s \models_i \phi \textit{ if and only if } E, s \models_i B_i\phi$$

This result states that agent i believes ϕ in state s if and only if ϕ is one of the facts that are perceived as true by that agent in that state (in fact, the "only if" side of the proposition is the classical axiom of knowledge, axiom T).

Results on Negative Introspection

Proposition 8 (Single-agent Negative Introspection)

$$E, s \models_i \neg B_i\phi \textit{ implies } E, s \models_i B_i\neg B_i\phi$$

This proposition states that axiom 5 (the classical axiom of negative introspection) holds for each belief operator B_i (*i.e.* every agent has introspective capabilities on its own negative beliefs).

Proposition 9 (Inter-agent Negative Introspection)

For two different agents i, j, it does not hold that

$$E, s \models_i \neg B_j\phi \textit{ implies } E, s \models_i B_j\neg B_j\phi$$

This result states that each agent i is aware of the fact that, even if it has reasons to think that j does not believe ϕ, it may just be the case that j believes ϕ indeed (and, therefore, j would believe that it believed ϕ). Thus, it is another expression of the uncertainty that an agent has about the beliefs of the other agents. However, the following fact does hold:

Proposition 10 (Generation of Negative Beliefs)

$$E, s \models_i \neg B_j\phi \textit{ implies } E, s \models_i B_i\neg B_j\phi$$

This proposition is expressing the fact that agent i can make positive introspection on negated beliefs of other agents. It is related to the following result, that characterizes an agent's set of propositional negative beliefs:

Proposition 11 (Characterization of Negative Beliefs)
For any propositional formula ϕ,

$$E, s \dashv_i \phi \textit{ if and only if } E, s \models_i \neg B_i\phi$$

Agent i does not believe ϕ at s if and only if ϕ is one the facts that is rejected by i at s.

Summarising the main results shown in this section:

- All forms of logical omniscience are avoided.
- Each agent is aware of its positive and negative beliefs, and is also aware of the fact that the other agents enjoy this introspective capability. However, an agent is uncertain about the way the present situation is perceived by other agents and, therefore, it is unable to know anything about the other agent's beliefs.
- The positive and negative beliefs of an agent in an state reflect, as our intuitions suggested, the facts that are taken as true or false by the agent in that state.

5 Comparison with Previous Proposals

The most outstanding difference of our proposal with previous works is the idea of considering *subjective* situations, that may be perceived in different ways by different agents. Technically, this fact implies two differences of our approach with respect to others:

- A situation is described with two functions ($\mathcal{T}_i$ and $\mathcal{F}_i$) *for each agent* i.
- Two satisfiability and unsatisfiability relations between situations and formulas ($\models_i$ and $\dashv_i$) are also defined *for each agent*.

The rest of the section is devoted to the comparison of our proposal with the two approaches to the problem of logical omniscience with which it shares more similarities: Levesque's *logic of explicit and implicit beliefs* ([5]) and Thijsse's *hybrid sieve systems* ([6]) (a more detailed review of these approaches may be found at [7]).

5.1 Levesque's Logic of Implicit and Explicit Beliefs

Levesque uses a language with two modal operators: B for *explicit* beliefs and L for *implicit* beliefs. These operators are not allowed to be nested in the formulas of the language. A *structure for explicit and implicit beliefs* is defined as a tuple M=(S, $\mathcal{B}$, T, F), where S is the set of primitive situations, $\mathcal{B}$ is a subset of S that represents the situations that could be the actual one and T and F are functions from the set of primitive propositions into subsets of S. Intuitively, $T(P)$ contains all the situations that support the truth of P, whereas $F(P)$ contains the ones that support the falsehood of P. A situation s can be *partial* (if there is a primitive proposition which is neither true nor false in s) and/or *incoherent* (if there is a proposition which is both true and false in s). A situation is *complete* if it is neither partial nor incoherent. A complete situation s is *compatible* with a situation t if s and t agree in all the points in which t is defined. $\mathcal{B}^*$ is the set of all complete situations of S that are compatible with some situation in $\mathcal{B}$.

The relations $\models_T$ and $\models_F$ between situations and formulas are defined as follows:

- $M,s \models_T P$, where P is a primitive proposition, if and only if $s \in T(P)$
- $M,s \models_F P$, where P is a primitive proposition, if and only if $s \in F(P)$
- $M,s \models_T \neg\varphi$ if and only if $M,s \models_F \varphi$
- $M,s \models_F \neg\varphi$ if and only if $M,s \models_T \varphi$
- $M,s \models_T (\varphi \wedge \psi)$ if and only if $M,s \models_T \varphi$ and $M,s \models_T \psi$
- $M,s \models_F (\varphi \wedge \psi)$ if and only if $M,s \models_F \varphi$ or $M,s \models_F \psi$
- $M,s \models_T B\varphi$ if and only if $M,t \models_T \varphi \;\; \forall t \in \mathcal{B}$
- $M,s \models_F B\varphi$ if and only if $M,s \not\models_T B\varphi$
- $M,s \models_T L\varphi$ if and only if $M,t \models_T \varphi \;\; \forall t \in \mathcal{B}^*$
- $M,s \models_F L\varphi$ if and only if $M,s \not\models_T L\varphi$

There are some similarities between our approach and Levesque's logic of implicit and explicit beliefs. However, they are more apparent than real, as shown in this listing:

- Levesque also considers a satisfiability and an unsatisfiability relation between situations and doxastic formulas.
  However, these relations are not considered *for each agent*.
- Levesque also describes each situation with two functions $\mathcal{T}$ and $\mathcal{F}$.
  Again, these functions are not indexed by each agent, as our functions are (Levesque considers an objective description of what is true and what is false in each situation). Another important difference is that Levesque's functions deal with basic propositions, and not with formulas as our functions do.
- Both approaches allow the presence of *partial* or *inconsistent* situations.
  However note that, in our case, it is not the (objective) description of the situation that is partial or inconsistent, but the *subjective* perception that an agent may have of it. Thus, the notions of partiality and inconsistency have a much more natural interpretation in our framework.
- Both approaches avoid all the forms of logical omniscience.
  The reason is different in each case, though. In Levesque's logic of explicit and implicit beliefs, it is the presence of incoherent situations that prevents logical omniscience. In our proposal, there is no need to have inconsistent situations to avoid logical omniscience. In fact, we solve that problem by defining $\mathcal{T}_i$ and $\mathcal{F}_i$ over arbitrary sets of formulas, and not over basic propositions.
- There are accessibility relations between situations for each agent in both systems.
  Levesque's accessibility relation between situations is left implicit; our accessibility relations are explicit. Furthermore, the intuition underlying these relations is somewhat different, as explained in section 2.

Many differences with Levesque's approach have just been outlined. Other differences that may be mentioned are the following:

- Levesque only considers one agent, and does not allow nested beliefs[6]. Thus, his agents do not have any introspective capabilities.
- Levesque defines *explicit* and *implicit* beliefs, whereas we do not make this distinction.
- Even though Levesque avoids logical omniscience, his agents must necessarily believe all those tautologies that are formed by *known* basic propositions (those propositions P for which the agent believes $(P \vee \neg P)$), regardless of their complexity. This is not the case in our approach, because we deal directly with formulas.
- There is a different treatment of the unsatisfiability relation when applied to beliefs, because he transforms $\dashv$ into $\not\models$, whereas we don't.

[6] However, similar logics that deal with many agents have been considered by Halpern and Lakemeyer, see *e.g.* [10].

5.2 Thijsse's Hybrid Sieve Systems

Thijsse ([6]) proposes a way of using *partial logics* to deal with various forms of logical omniscience. He defines a *partial model* as a tuple $(W, \mathcal{B}_1, \ldots, \mathcal{B}_n, V)$, where W is a set of worlds, $\mathcal{B}_i$ is the accessibility relation between worlds for agent i and V is a *partial* truth assignment to the basic propositions in each world. $\top$ is a primitive proposition that is always interpreted as *true*. Truth ($\models$) and falsity ($\Dashv$) relations are defined in the following way:

- $M,w \models \top$
- $M,w \not\Dashv \top$
- $M,w \models P$, where P is a primitive proposition, iff $V(P,w)= 1$
- $M,w \Dashv P$, where P is a primitive proposition, iff $V(P,w)= 0$
- $M,w \models \neg\varphi$ iff $M,w \Dashv \varphi$
- $M,w \Dashv \neg\varphi$ iff $M,w \models \varphi$
- $M,w \models (\varphi \wedge \psi)$ iff $M,w \models \varphi$ and $M,w \models \psi$
- $M,w \Dashv (\varphi \wedge \psi)$ iff $M,w \Dashv \varphi$ or $M,w \Dashv \psi$
- $M,w \models B_i\varphi$ iff $M,v \models \varphi$ $\forall v$ such that (w, v) $\epsilon\mathcal{B}_i$
- $M,w \Dashv B_i\varphi$ iff $\exists v$ s.t. (w, v) $\epsilon\mathcal{B}_i$ and $M,v \Dashv \varphi$

The most important similarities between our approach and Thijsse's are the following:

- n agents and n explicit accessibility relations are considered.
  However, as in Levesque's case, there are no restrictions on these relations, and the intuitive meaning of our accessibility relations is slightly different.
- Two relations (of satisfiability and unsatisfiability) are defined. Moreover, a similar clause is used to provide a meaning to the unsatisfiability relation with respect to the belief operator.
  As before, the main difference is that we provide two relations *for each agent*.
- There are no tautologies in Thijsse's system; therefore, he does not have to care about some forms of logical omniscience (closure under valid implication and belief of valid formulas).
- Closure under material implication and closure under conjunction do not hold in Thijsse's approach either.

The main difference with Thijsse's proposal is that he uses *partial* assignments of truth values *over basic propositions* for each state; thus, a proposition may be true, false or undefined in each state. We deal with formulas, not with basic propositions, and each formula may be supported *and/or* rejected by *each agent* in each state. Therefore, Thijsse's approach is three-valued, whereas ours is more of a four-valued kind, such as Levesque's.

6 Summary

It has been argued that each agent perceives its actual situation in a particular way, which may be different from that of other agents located in the same situation. This intuitive idea has been formalized with the notion of *subjective situations*. These entities are the base of a doxastic logic, in which the meaning of the belief operators seems to fit with the general intuitions about how the doxastic attitude of a non-ideal agent should behave. In particular, logical omniscience is avoided while some interesting introspective properties are maintained.

References

1. R. Fagin, J. Halpern, Y. Moses, M. Vardi, *Reasoning about knowledge*, MIT Press, 1995.
2. J. Hintikka, *Knowledge and belief*, Cornell University Press, Ithaca, N.Y., 1962.
3. S. Kripke, A semantical analysis of modal logic I: normal modal propositional calculi, *Zeitschrift für Mathematische Logik und Grundlagen Mathematik* **9** (1963), 67-96.
4. J. Hintikka, Impossible possible worlds vindicated, *Journal of Philosophical Logic* **4** (1975), 475-484.
5. H.J. Levesque, A logic of implicit and explicit belief, *Proceedings of the Conference of the American Association for Artificial Intelligence, AAAI-84* (1984), 198-202.
6. E. Thijsse, Combining partial and classical semantics. A hybrid approach to belief and awareness, in P. Doherty (Ed.), *Partiality, modality, and nonmonotonicity*, Studies in Logic, Language and Information, Center for the Study of Language and Information Publications (1996), 223-249.
7. A. Moreno, Avoiding logical omniscience and perfect reasoning: a survey, *AI Communications* **2** (1998), 101-122.
8. A. Moreno, T. Sales, Dynamic belief analysis, in J. Müller, M. Wooldridge, N. Jennings (Eds.), *Intelligent Agents III: Agent theories, architectures and languages*, Lecture Notes in Artificial Intelligence 1193, Springer Verlag (1997), 87-102.
9. A. Moreno, T. Sales, Limited logical belief analysis, in L. Cavedon, A. Rao, W. Wobcke (Eds.), *Intelligent Agent Systems: Theoretical and Practical Issues*, Lecture Notes in Artificial Intelligence 1209, Springer Verlag (1997), 104-118.
10. J. Halpern, G. Lakemeyer, Multi-agent only knowing, in *Proceedings of the Sixth Conference on Theoretical Aspects of Rationality and Knowledge, TARK-96* (1996), 251-265.

This article was processed using the LaTeX macro package with LLNCS style