

Dynamic Epistemic Logic for Implicit and Explicit Beliefs

Fernando R. Velázquez-Quesada

Abstract

The *dynamic* turn in Epistemic Logic is based on the idea that notions of information should be studied together with the actions that modify them. Dynamic epistemic logics have explored how knowledge and beliefs change as consequence of, among others, acts of observation and upgrade. Nevertheless, the omniscient nature of the represented agents has kept finer actions outside the picture, the most important being the action of *inference*.

Following proposals for representing non-omniscient agents, recent works have explored how *implicit* and *explicit knowledge* change as a consequence of acts of observation, inference, consideration and even forgetting. The present work proposes a further step towards a common framework for representing finer notions of information and their dynamics. We propose a combination of existing works in order to represent *implicit* and *explicit beliefs*. Then, after adapting definitions for the actions of *upgrade* and *retraction*, we discuss the action of *inference on beliefs*, analyzing its differences with respect to inference on knowledge and proposing a rich system for its representation.

1 Introduction

Epistemic Logic [21] and its possible worlds semantics is a powerful and compact framework for representing an agent's information. Their *dynamic* versions [14] have emerged to analyze not only information in its knowledge and belief versions, but also the actions that modify them. Nevertheless, agents represented in this framework are *logically omniscient*, that is, their information is closed under logical consequence. This property, useful in some applications, hides finer reasoning actions that are crucial in some others, the most important being that of *inference*.

Based on the awareness approach of [17], several works have explored dynamics of information for non-omniscient agents. In a *propositional dynamic logic (PDL)* style, some of them have explored how the act of inference modifies an agent's *explicit* knowledge [15; 22]. In a *dynamic epistemic* style, some others have explored how the acts of observation, inference, consideration and forgetting affect *implicit* and *explicit knowledge* [5; 18; 10; 13].

The present work follows the previous ones, now focussing on the notion of *beliefs*. We combine approaches of the existing literature, proposing a setting for representing the notions of *implicit* and *explicit belief* (Section 2). Then we look into the dynamics of these notions; first, by adapting existing proposals to define the actions of *explicit upgrade* (*explicit revision*) and *retraction* (Section 3), and second, by discussing the action of *inference on beliefs* and its differences with *inference on knowledge*, and by proposing a rich system for its representation (Section 4).

2 Modelling implicit and explicit beliefs

This section recalls a framework for implicit and explicit information and a framework for beliefs. By combining them, we will get our model for representing implicit/explicit beliefs. But before going into their details, we recall the framework on which all the others are based.

Epistemic Logic. The frameworks of this section are based on that of *Epistemic Logic* (EL; [21]). Given a set of atomic propositions P , the EL language extends the propositional one with formulas of the form $\Box \varphi$: “*the agent is informed about φ* ”. Though there are several possibilities, the classical semantic model for EL-formulas are *Kripke models*, tuples $M = \langle W, R, V \rangle$ with W a non-empty set of *possible worlds*, $V : W \rightarrow \wp(P)$ an atomic valuation function indicating which atomic propositions are true at each world, and $R \subseteq (W \times W)$ an accessibility relation indicating which worlds the agent considers possible from each one of them.

Formulas are evaluated on *pointed models* (M, w) with M a Kripke model and $w \in W$ a given evaluation point. Boolean connectives are interpreted as usual; the key clause is the one for $\Box \varphi$, indicating that *the agent is informed about φ at w iff φ is true in all the worlds the agent considers possible from w* :

$$(M, w) \models \Box \varphi \quad \text{iff} \quad \text{for all } u \in W, Rwu \text{ implies } (M, u) \models \varphi$$

2.1 Implicit and explicit information

Non-omniscient agents. The formula $\Box(\varphi \rightarrow \psi) \rightarrow (\Box \varphi \rightarrow \Box \psi)$ is valid in Kripke models: the agent’s information is closed under logical consequence. This becomes obvious when we realize that each possible world stands for a maximally consistent set of formulas. So if both $\Box(\varphi \rightarrow \psi)$ and $\Box \varphi$ hold at world w , both $\varphi \rightarrow \psi$ and φ are true in all worlds R -reachable from w . But then ψ also holds in all such worlds, and therefore $\Box \psi$ holds at w . Usually the discussion revolves around whether this is a reasonable assumption for ‘real’ agents. Even computational agents may not have this property, since they may lack of resources (space and/or time) to derive all the logical consequences of their information [1].

One of the most influential solutions to this *omniscience problem* is *awareness logic* [17]. This approach follows the idea of making a difference between *implicit (potential) information*, what the agent can eventually get, and *explicit information*, what the agent actually has [23; 25; 24; 27]. The main observation is that, in order to have explicit information about some formula φ , besides having it as implicit information, the agent should be *aware* of φ .

Syntactically, awareness logic extends the *EL* language with formulas of the form $A\varphi$: “the agent is aware of φ ”. Semantically, it extends Kripke models with a function A that assigns a set of formulas to the agent in each possible world. The new formulas are evaluated in the following way:

$$(M, w) \Vdash A\varphi \quad \text{iff} \quad \varphi \in A(w)$$

Implicit information about φ is defined as $\Box\varphi$, but explicit information is defined as $\Box\varphi \wedge A\varphi$. Although implicit information is still closed under logical consequence, explicit information is not. This follows from the fact that, different from the possible worlds, the A -sets *do not need* to have any closure property; in particular, $\{\varphi \rightarrow \psi, \varphi\} \subseteq A(w)$ does not imply $\psi \in A(w)$.

Agents with reasoning abilities. Still, though a ‘real’ agent’s information does not need to be closed under logical consequence, it does not need to be static. The more interesting approach for us is that in which the agent can extend her explicit information by the adequate actions. But, which are these actions and what does the agent needs in order to perform them?

In [15], the author proposes a framework in which the actions available to the agent are different rules (e.g., modus ponens, conjunction elimination), each one of them represented by a relation between worlds that should be faithful to the rule’s spirit (e.g., the modus ponens relation should connect worlds with an implication and its antecedent with worlds augmented with the consequent). This yields an agent that does not need to be omniscient, but still is able to perform inferences.

In [22] the author goes one step further: a rule cannot be used by an agent unless the rule itself is also part of her explicit information. For example, for two worlds to be connected by the modus ponens relation, the initial one should have not only an implication and its antecedent, but also the *modus ponens* rule itself.

The combination of the mentioned ideas have produced models for representing implicit and explicit *knowledge* ([5; 28; 13; 18; 10] among others). But the notion of *belief* is different, as we discuss in the next subsection.

2.2 Modelling beliefs

The *KD45* approach. For modelling knowledge in *EL*, it is usually asked for the accessibility relation R to be at least *reflexive* (making $\Box\varphi \rightarrow \varphi$ valid:

if the agent knows φ , then φ is true), and often to be also *transitive* and *euclidean* (giving the agent full positive and negative introspection). Beliefs can be represented in a similar way, now asking for R to satisfy weaker properties, the crucial one following the idea that, though beliefs do not need to be true, we can expect them to be *consistent*. This is achieved by asking for the relation to be *serial*, making the D axiom $\neg\Box\perp$ valid. Full introspection is usually assumed, yielding the classical $KD45$ approach.

Belief as what is most plausible. But beliefs are different from knowledge. Intuitively, we do not believe something because it is true in all possible situations; we believe it because it is true in those we consider most likely to be the case [19; 26]. This idea has led the development of variants of Kripke models [12; 4; 3]. Here we recall the *plausibility models* of [3].

A *plausibility model* is a Kripke model in which the accessibility relation, denoted now by \leq , is interpreted as a *plausibility* relation ordering possible worlds. This relation is assumed to be a preorder (a reflexive and transitive relation). Moreover, since the idea is to define the agent's beliefs as what is true in the most plausible worlds from the evaluation point, \leq should satisfy an important extra property: for any possible world w , the set of worlds that are better than w among those comparable to it should have maximal worlds. In order to state this property formally, denote by V_w the set of worlds comparable to w (its comparability class: $V_w := \{u \mid w \leq u \text{ or } u \leq w\}$) and by $\text{Max}_{\leq}(U)$ the set of \leq -maximal worlds of U ($\text{Max}_{\leq}(U) := \{v \in U \mid \text{for all } u \in U, u \leq v\}$). Then, in a plausibility model, the accessibility relation \leq is asked to be a *locally well-preorder*: a reflexive and transitive relation such that, for each comparability class V_w and for every non-empty $U \subseteq V_w$, $\text{Max}_{\leq}(U) \neq \emptyset$. Note how the existence of maximal elements in every $U \subseteq V_w$ implies the already required reflexivity, but also *connectedness* inside V_w . In particular, if two worlds w_2 and w_3 are more plausible than a given w_1 ($w_1 \leq w_2$ and $w_1 \leq w_3$), then these two worlds should be \leq -related ($w_2 \leq w_3$ or $w_3 \leq w_2$ or both).

Interestingly, the agent's indistinguishability relation can be derived from the plausibility one. If two worlds are \leq -related, then though the agent considers one of them more plausible than the other, she cannot discard one of them when the other one is given. In other words, worlds that are \leq related are in fact epistemically indistinguishable.

For the language we have two options.¹ We can extend the propositional language with formulas of the form $B\varphi$, semantically interpreted as

$$(M, w) \Vdash B\varphi \quad \text{iff} \quad \text{for all } u \in W, u \in \text{Max}_{\leq}(R_{\leq}(w)) \text{ implies } (M, u) \Vdash \varphi, \\ \text{where } R_{\leq}(w) := \{u \in W \mid w \leq u\}.$$

¹In fact, the mentioned works, [12; 4; 3], use the notion of *conditional* belief as the primitive one, rather than plain belief. We have chosen to stick with the notion of plain belief through the present notes, leaving an analysis of the notions of implicit/explicit conditional beliefs for further work.

The second option is to use a standard modal language with $[\leq]$ standing for the relation \leq , and then define beliefs in terms of it. Given the properties of \leq (in particular, reflexivity, transitivity and connectedness), it is not hard to see that φ is true in the most plausible worlds from w iff w can see a better world from which all successors are φ worlds. This yields the following definition for “the agent believes φ ”:

$$B\varphi := \langle \leq \rangle [\leq] \varphi$$

2.3 Combining the models

Our framework for representing implicit and explicit beliefs combines the mentioned ideas. The language has two components: formulas and rules. Formulas are given by a propositional language extended, first, with modalities $\langle \leq \rangle$ and $\langle \simeq \rangle$, and second, with formulas of the form $A \varphi$ and $R \rho$, where φ is a formula and ρ a rule. Rules, on the other hand, are pairs consisting of a set of formulas, the rule’s premises, and a single formula, the rule’s conclusion. The formal definition of our language is as follows.

Definition 2.1 (Language \mathcal{L}). Given a set of atomic propositions P , formulas φ and rules ρ of the *plausibility-access* language \mathcal{L} are given, respectively, by

$$\begin{aligned} \varphi &::= p \mid A \varphi \mid R \rho \mid \neg \varphi \mid \varphi \vee \psi \mid \langle \simeq \rangle \varphi \mid \langle \leq \rangle \varphi \\ \rho &::= (\{\varphi_1, \dots, \varphi_{n_\rho}\}, \psi) \end{aligned}$$

where $p \in P$. Formulas of the form $A \varphi$ are read as “the agent has acknowledged that formula φ is true”, and formulas of the form $R \rho$ as “the agent has acknowledged that rule ρ is truth-preserving”. For the modalities, $\langle \leq \rangle \varphi$ is read as “there is a more plausible world where φ holds”, and $\langle \simeq \rangle \varphi$ as “there is an epistemically indistinguishable world where φ holds”. Other boolean connectives as well as the box modalities $[\simeq]$ and $[\leq]$ are defined as usual. We denote by \mathcal{L}_f the set of formulas of \mathcal{L} , and by \mathcal{L}_r its set of rules.

Though rules are usually presented as schemas, our rules are defined as particular instantiations (e.g., the rule $(\{p \wedge q\}, p)$ is different from the rule $(\{q \wedge r\}, q)$). Since they will be applied in a generalized modus ponens form (if the agent has all the premises, she can derive the conclusion), using concrete formulas avoids details of instantiation, therefore facilitating the definition. When dealing with them, the following definitions will be useful.

Definition 2.2. Given a rule ρ , we will denote its *set of premises* by $\text{pm}(\rho)$, its *conclusion* by $\text{cn}(\rho)$, and its *translation* (an implication whose antecedent is the finite conjunction of ρ ’s premises and whose consequent is ρ ’s conclusion) by $\text{tr}(\rho)$.

For the semantic model, we will extend the described plausibility models with two functions.

Definition 2.3 (Plausibility-access model). With P the set of atomic propositions, a *plausibility-access (PA) model* is a tuple $M = \langle W, \leq, V, A, R \rangle$ where $\langle W, \leq, V \rangle$ is a plausibility model over P and

- $A : W \rightarrow \wp(\mathcal{L}_f)$ is the *access set function*, assigning to the agent a set of formulas of \mathcal{L} in each possible world,
- $R : W \rightarrow \wp(\mathcal{L}_r)$ is the *rule set function*, assigning to the agent a set of rules of \mathcal{L} in each possible world.

Functions A and R can be seen as valuations with a particular range, assigning to the agent a set of formulas and a set of rules at each possible world, respectively. Moreover, recall that two worlds that are \leq related are epistemically indistinguishable, so we define \simeq as the union of \leq and its converse ($\simeq := \leq \cup \geq$): the agent cannot distinguish between two worlds if she considers one of them more plausible than the other.

A *pointed plausibility-access model* (M, w) is a plausibility-access model with a distinguished world $w \in W$.

Here it is important to emphasize our interpretation of the A -sets. Different from [17] and [10], we do not interpret them as “the formulas the agent is aware of at world w ”, but rather as “the formulas the agent has acknowledged as true at world w ”, closer to the ideas in [15; 22; 18].

Now the semantic evaluation. The modalities $\langle \leq \rangle$ and $\langle \simeq \rangle$ are interpreted via their correspondent relation in the usual way, and formulas of the form $A\varphi$ and $R\rho$ are interpreted with our two new functions.

Definition 2.4 (Semantic interpretation). Let (M, w) be a pointed PA model with $M = \langle W, \leq, V, A, R \rangle$. Atomic propositions and boolean operators are interpreted as usual. For the remaining cases,

$$\begin{aligned}
(M, w) \Vdash A\varphi & \quad \text{iff} \quad \varphi \in A(w) \\
(M, w) \Vdash R\rho & \quad \text{iff} \quad \rho \in R(w) \\
(M, w) \Vdash \langle \leq \rangle \varphi & \quad \text{iff} \quad \text{there is a } u \in W \text{ such that } w \leq u \text{ and } (M, u) \Vdash \varphi \\
(M, w) \Vdash \langle \simeq \rangle \varphi & \quad \text{iff} \quad \text{there is a } u \in W \text{ such that } w \simeq u \text{ and } (M, u) \Vdash \varphi
\end{aligned}$$

For characterizing valid formulas, an important observation is that a locally well-preorder is a *locally connected and conversely well-founded pre-order* [3]. Then, by standard results on canonicity and modal correspondence (Chapter 4 of [11]), the axiom system of Section 2.6 of [3] (Table 1) is also sound and (weakly) complete for our language \mathcal{L} with respect to ‘non-standard’ plausibility-access models: those in which \leq is reflexive, transitive and locally connected (axioms T_{\leq} , 4_{\leq} and LC) and \simeq the symmetric extension of \leq (axioms T_{\simeq} , 4_{\simeq} , B_{\simeq} and Inc). But such models also have the *finite model property* (with respect to formulas in our language), so completeness

<i>Prop</i>	$\vdash \varphi$ for φ a propositional tautology	<i>MP</i>	If $\vdash \varphi \rightarrow \psi$ and $\vdash \varphi$, then $\vdash \psi$
K_{\leq}	$\vdash [\leq](\varphi \rightarrow \psi) \rightarrow ([\leq]\varphi \rightarrow [\leq]\psi)$	K_{\simeq}	$\vdash [\simeq](\varphi \rightarrow \psi) \rightarrow ([\simeq]\varphi \rightarrow [\simeq]\psi)$
<i>Dual</i> _{\leq}	$\vdash \langle \leq \rangle \varphi \leftrightarrow \neg[\leq]\neg\varphi$	<i>Dual</i> _{\simeq}	$\vdash \langle \simeq \rangle \varphi \leftrightarrow \neg[\simeq]\neg\varphi$
<i>Nec</i> _{\leq}	If $\vdash \varphi$, then $\vdash [\leq]\varphi$	<i>Nec</i> _{\simeq}	If $\vdash \varphi$, then $\vdash [\simeq]\varphi$
T_{\leq}	$\vdash [\leq]\varphi \rightarrow \varphi$	T_{\simeq}	$\vdash [\simeq]\varphi \rightarrow \varphi$
4_{\leq}	$\vdash [\leq]\varphi \rightarrow [\leq][\leq]\varphi$	4_{\simeq}	$\vdash [\simeq]\varphi \rightarrow [\simeq][\simeq]\varphi$
		B_{\simeq}	$\vdash \varphi \rightarrow [\simeq]\langle \simeq \rangle \varphi$
<i>LC</i>	$(\langle \simeq \rangle \varphi \wedge \langle \simeq \rangle \psi) \rightarrow (\langle \simeq \rangle (\varphi \wedge \langle \leq \rangle \psi) \vee \langle \simeq \rangle (\psi \wedge \langle \leq \rangle \varphi))$		
<i>Inc</i>	$\langle \leq \rangle \varphi \rightarrow \langle \simeq \rangle \varphi$		

Table 1: Axiom system for \mathcal{L} with respect to plausibility-access models.

with respect to plausibility-access models follows from the fact that every strict preorder is conversely well-founded.

Note how the axiom system does not have axioms for formulas of the form $A\varphi$ and $R\rho$. This is because, as mentioned before, such formulas are simply special *atoms* for the dedicated *valuation functions* A and R . Moreover, we have not asked for the A - and R -sets to have any special closure property and there is no restriction in the way they interact with each other.² Just like axiom systems for Epistemic Logic do not require special axioms describing the behaviour of atomic propositions (unless, of course, they have special properties, like q being true every time p is, characterized by $p \rightarrow q$), our system does not require special axioms for these special atoms. More precisely, in the canonical model construction, we only need to define access and rule sets in the proper way:

$$A(w) := \{\varphi \in \mathcal{L}_f \mid A\varphi \in w\} \quad R(w) := \{\rho \in \mathcal{L}_r \mid R\rho \in w\}$$

Then, formulas of the form $A\varphi$ and $R\rho$ also satisfy the crucial *Truth Lemma*, and completeness follows. Again, see Chapter 4 of [11] for details.

2.3.1 Implicit and explicit beliefs

It is time to define the notions of implicit and explicit beliefs. Our definitions, shown in Table 2, combine ideas from [3], [18] and [10]. Note how the agent believes the formula φ (the rule ρ) implicitly iff φ ($\text{tr}(\rho)$) is true in the most plausible worlds, but in order to believe it explicitly, the agent should also acknowledge φ (ρ) as true (truth-preserving) in these ‘best’ worlds.

Explicit beliefs are implicit beliefs, witness the following validities:

$$B_{\text{Ex}}\varphi \rightarrow B_{\text{Im}}\varphi \quad B_{\text{Ex}}\rho \rightarrow B_{\text{Im}}\rho$$

²In [17], the authors explore and characterize several closure properties of A -sets.

The agent <i>implicitly</i> believes formula φ	$B_{\text{Im}}\varphi := \langle \leq \rangle [\leq] \varphi$
The agent <i>explicitly</i> believes formula φ	$B_{\text{Ex}}\varphi := \langle \leq \rangle [\leq] (\varphi \wedge \mathbf{A} \varphi)$
The agent <i>implicitly</i> believes rule ρ	$B_{\text{Im}}\rho := \langle \leq \rangle [\leq] \text{tr}(\rho)$
The agent <i>explicitly</i> believes rule ρ	$B_{\text{Ex}}\rho := \langle \leq \rangle [\leq] (\text{tr}(\rho) \wedge \mathbf{R} \rho)$

Table 2: Implicit and explicit beliefs about formulas and rules.

A possibly more interesting point is the following. An agent in [17; 10] is non-omniscient due to lack of attention; she does not need to be *aware of* every formula. On the other hand, our agent is aware of all formulas, but still she is non-omniscient because she does not need to be *aware that* a formula is true. This may seem a small difference, but the interpretation of the \mathbf{A} sets determines the reasonable operations over them. An agent can become aware of any formula at any time, so any formula can be added to the \mathbf{A} -sets without further requirement [10]. On the other hand, it is a stretch to assume that an agent can recognize as true any formula at any moment; it is more reasonable to ask for some derivation device, which in this work will be a *rule application* [15; 22; 18].

We finish this section by mentioning some properties of implicit and explicit beliefs about formulas. (Rules behave in a similar way.) Implicit beliefs are closed under logical consequence: if the most plausible worlds satisfy both $\varphi \rightarrow \psi$ and φ , then they also satisfy ψ . But explicit beliefs do not need to have this property because the \mathbf{A} -sets do not need to have any closure property: having φ and $\varphi \rightarrow \psi$ does not guarantee to have ψ .

Though \leq is reflexive, neither implicit nor explicit beliefs have to be true because the real world does not need to be among the most plausible ones. Nevertheless, reflexivity makes implicit (and therefore explicit) beliefs consistent. Every world has at least one \leq -successor, so $\neg B_{\text{Im}}\perp$ is valid.

Implicit beliefs are positively and negatively introspective. This is the case because the notion of ‘most plausible worlds’ is global inside the same comparability class. For positive introspection, if the set of maximal worlds contains only φ -worlds ($B_{\text{Im}}\varphi$), so does the maximal from the maximal ones ($B_{\text{Im}}B_{\text{Im}}\varphi$). And for negative introspection, if there is a $\neg\varphi$ -world u in the maximal worlds, ($\neg B_{\text{Im}}\varphi$), then u is also in those maximal from the maximal ones ($B_{\text{Im}}\neg B_{\text{Im}}\varphi$). But this does not extend to explicit beliefs, again because the \mathbf{A} -sets do not need to have any closure property. Having φ does not guarantee to have $B_{\text{Ex}}\varphi$ (so $B_{\text{Ex}}\varphi \rightarrow B_{\text{Ex}}B_{\text{Ex}}\varphi$ is not valid), and *not* having φ does not guarantee to have $\neg B_{\text{Ex}}\varphi$ (so $\neg B_{\text{Ex}}\varphi \rightarrow B_{\text{Ex}}\neg B_{\text{Ex}}\varphi$ is not valid).

3 Dynamics part one: upgrade and retraction

We have a framework for representing implicit and explicit beliefs. We now look at their *dynamics* by introducing two actions that modify them.

3.1 Explicit upgrade

The χ -upgrade operation [4; 3] modifies the plausibility relation \leq to put the χ -worlds at the top, therefore *revising* the agent's beliefs. Here we have two possibilities, depending on whether it also adds χ to the \mathbf{A} -sets (*explicit upgrade*) or not (*implicit upgrade*). Here is the definition of the first case.

Definition 3.1 (Explicit upgrade). Let $M = \langle W, \leq, V, \mathbf{A}, \mathbf{R} \rangle$ be a PA model and χ a formula in \mathcal{L} . The PA model $M_{\chi\uparrow^+} = \langle W, \leq', V, \mathbf{A}', \mathbf{R} \rangle$ differs from M in the plausibility relation and in the access set function:

$$\begin{aligned} \leq' &:= (\leq; \chi?) \cup (\neg\chi?; \leq) \cup (\neg\chi?; \simeq; \chi?) \\ \mathbf{A}'(w) &:= \mathbf{A}(w) \cup \{\chi\} \quad \text{for every } w \in W \end{aligned}$$

Note how the *upgrade* operation is functional: for every model M it returns one and only one model $M_{\chi\uparrow^+}$.

The new plausibility relation is given in a PDL style: we have $w \leq' u$ iff (1) $w \leq u$ and u is a χ -world, or (2) w is a $\neg\chi$ -world and $w \leq u$, or (3) $w \simeq u$, w is a $\neg\chi$ -world and u is a χ -world. There are other possible definitions for \leq' [4; 3], and the chosen one, so-called *radical upgrade*, is just an example of what can be defined.

The operation preserves models in the intended class.

Proposition 1. *If M is a PA model, then so is $M_{\chi\uparrow^+}$.*

We extend the language to express the effect of an explicit upgrade; formulas of the form $\langle \chi \uparrow^+ \rangle \varphi$ are read as “it is possible to perform an explicit χ -upgrade after which φ holds”. There is no precondition for this action (the agent can perform an explicit upgrade whenever she wants), so the semantic interpretation is as follows.

Definition 3.2. Let (M, w) be a pointed PA model:

$$(M, w) \Vdash \langle \chi \uparrow^+ \rangle \varphi \quad \text{iff} \quad (M_{\chi\uparrow^+}, w) \Vdash \varphi$$

Note how the operation puts on top those worlds that are χ -worlds in the original model, but they do not need to be χ -ones after the upgrade. The plausibility relation changes, therefore changing the truth-value of formulas containing the modalities for \leq and/or \simeq and, in particular, changing the agent's beliefs. This is not strange at all, and in fact it corresponds to the

well-known Moore-like sentences (“ p is the case and you do not know it”) in *Public Announcement Logic* that become false after being announced, and therefore cannot be known.

Nevertheless, the operation behaves as expected for propositional formulas. The operation does not change valuations, so if χ is purely propositional, the operation will put current χ -worlds on top, and they will still be χ -worlds after the operation so the agent will believe χ .

The validities in our new language can be axiomatized by using *reduction axioms*, valid formulas that indicate how to translate a formula with the new modality $\langle \chi \uparrow^+ \rangle$ into a provably equivalent one without them. Then, completeness follows from the completeness of the basic system. We refer to [8] for an extensive explanation of this technique.

Theorem 1. *The axiom system of Table 1 together with axioms and rules of Table 3 (with \top the always true formula) provide a sound and (weakly) complete axiom system for formulas in the language \mathcal{L} plus the explicit upgrade modality with respect to plausibility-access models.*

$\vdash \langle \chi \uparrow^+ \rangle p \leftrightarrow p$	$\vdash \langle \chi \uparrow^+ \rangle A\chi \leftrightarrow \top$
$\vdash \langle \chi \uparrow^+ \rangle \neg\varphi \leftrightarrow \neg\langle \chi \uparrow^+ \rangle \varphi$	$\vdash \langle \chi \uparrow^+ \rangle A\varphi \leftrightarrow A\varphi \quad \text{for } \varphi \neq \chi$
$\vdash \langle \chi \uparrow^+ \rangle (\varphi \vee \psi) \leftrightarrow (\langle \chi \uparrow^+ \rangle \varphi \vee \langle \chi \uparrow^+ \rangle \psi)$	$\vdash \langle \chi \uparrow^+ \rangle R\rho \leftrightarrow R\rho$
$\vdash \langle \chi \uparrow^+ \rangle \langle \leq \rangle \varphi \leftrightarrow \langle \leq \rangle (\chi \wedge \langle \chi \uparrow^+ \rangle \varphi) \vee (\neg\chi \wedge \langle \leq \rangle \langle \chi \uparrow^+ \rangle \varphi) \vee (\neg\chi \wedge \langle \simeq \rangle (\chi \wedge \langle \chi \uparrow^+ \rangle \varphi))$	
$\vdash \langle \chi \uparrow^+ \rangle \langle \simeq \rangle \varphi \leftrightarrow \langle \simeq \rangle \langle \chi \uparrow^+ \rangle \varphi$	
From $\vdash \varphi$ infer $\vdash \langle \chi \uparrow^+ \rangle \varphi$	

Table 3: Axioms and rules for explicit upgrade formulas.

The reduction axioms simply indicate how each kind of formula is affected by the explicit upgrade operation. For example, $\langle \chi \uparrow^+ \rangle p \leftrightarrow p$ states that atomic propositions are not affected, and both $\langle \chi \uparrow^+ \rangle A\chi \leftrightarrow \top$ and $\langle \chi \uparrow^+ \rangle A\varphi \leftrightarrow A\varphi$ for $\varphi \neq \chi$ together state that χ and only χ is added to the A-sets. The interesting axiom is the one for the plausibility modality $\langle \leq \rangle$. It is obtained with techniques from [9], and simply translates the three-cases PDL definition of the new plausibility relation: after an upgrade with χ there is a \leq -reachable world where φ holds iff before the operation (1) there is a \leq -reachable χ -world that will become φ after the upgrade, or (2) the current is a $\neg\chi$ -world that can \leq -reach another that will turn into a φ -one after the operation, or (3) the current is a $\neg\chi$ -world that can \simeq -reach another that is χ and will become φ after the upgrade. Similar reduction axioms have been presented in [9] in the context of preference upgrade.

3.2 Retraction

But there are also situations in which the agent simply retracts some explicit belief, that is, she does not acknowledge it as true anymore. This is achieved simply by removing the formula from the **A**-sets.

Definition 3.3 (Retraction). Let $M = \langle W, \leq, V, \mathbf{A}, \mathbf{R} \rangle$ be a PA model and χ a formula in \mathcal{L} . The PA model $M_{-\chi} = \langle W, \leq, V, \mathbf{A}', \mathbf{R} \rangle$ differs from M just in the access set function, given for every $w \in W$ as

$$\mathbf{A}'(w) := \mathbf{A}(w) \setminus \{\chi\}$$

Again, the *retraction* operation is functional. Moreover, it does not modify \leq , so it preserves plausibility models.

This operation is represented in the language by formulas of the form $\langle -\chi \rangle \varphi$, read as “it is possible to retract χ and after it φ holds”. Just like an upgrade, no precondition is needed.

Definition 3.4. Let (M, w) be a pointed PA model:

$$(M, w) \Vdash \langle -\chi \rangle \varphi \quad \text{iff} \quad (M_{-\chi}, w) \Vdash \varphi$$

Our definition behaves as we intend, witness the validity of $\langle -\chi \rangle \neg B_{\text{Ex}} \chi$: after retracting χ , the agent will not believe it explicitly.

For an axiom system, we can use reduction axioms again. In this case the axioms are simpler since only **A**-sets are affected.

Theorem 2. *The axiom system of Table 1 together with axioms and rules of Table 4 (with \perp the always false formula) provide a sound and (weakly) complete axiom system for formulas in the language \mathcal{L} plus the retraction modality with respect to plausibility-access models.*

$\vdash \langle -\chi \rangle p \leftrightarrow p$ $\vdash \langle -\chi \rangle \neg \varphi \leftrightarrow \neg \langle -\chi \rangle \varphi$ $\vdash \langle -\chi \rangle (\varphi \vee \psi) \leftrightarrow (\langle -\chi \rangle \varphi \vee \langle -\chi \rangle \psi)$ $\vdash \langle -\chi \rangle \langle \leq \rangle \varphi \leftrightarrow \langle \leq \rangle \langle -\chi \rangle \varphi$ $\vdash \langle -\chi \rangle \langle \simeq \rangle \varphi \leftrightarrow \langle \simeq \rangle \langle -\chi \rangle \varphi$ $\text{From } \vdash \varphi \text{ infer } \vdash \langle -\chi \rangle \varphi$	$\vdash \langle -\chi \rangle \mathbf{A} \chi \leftrightarrow \perp$ $\vdash \langle -\chi \rangle \mathbf{A} \varphi \leftrightarrow \mathbf{A} \varphi \quad \text{for } \varphi \neq \chi$ $\vdash \langle -\chi \rangle \mathbf{R} \rho \leftrightarrow \mathbf{R} \rho$
---	--

Table 4: Axioms and rules for retraction formulas.

Here the key axioms are $\langle -\chi \rangle \mathbf{A} \chi \leftrightarrow \perp$ and $\langle -\chi \rangle \mathbf{A} \varphi \leftrightarrow \mathbf{A} \varphi$ for $\varphi \neq \chi$, stating that χ and only χ is removed from the **A**-sets. Similar reduction axioms have been presented in [10] in the context of dynamics of awareness.

4 Dynamics part two: inference on beliefs

We now turn into the main part of our work. In this section we analyze *rule-based inference on beliefs*. We start by recalling the case of rule-based inference on *knowledge* [18].

The definition of implicit and explicit knowledge are simpler than those for beliefs since they depend directly on all the worlds the agent considers possible. The agent knows φ implicitly when it holds in all the worlds she considers possible, and she knows φ explicitly when she also recognizes it as true in all such worlds. The definitions of explicit knowledge about a rule ρ is given in a similar way.

$$\begin{aligned} K_{\text{Im}}\varphi &:= [\simeq]\varphi & K_{\text{Ex}}\varphi &:= [\simeq](\varphi \wedge A\varphi) \\ K_{\text{Im}}\rho &:= [\simeq]\text{tr}(\rho) & K_{\text{Ex}}\rho &:= [\simeq](\text{tr}(\rho) \wedge R\rho) \end{aligned}$$

The action of *inference on knowledge* with rule σ is defined as an operation that adds σ 's conclusion to the A -set of those worlds where the agent knows explicitly σ and its premises ($K_{\text{Ex}}\sigma \wedge K_{\text{Ex}}\text{pm}(\sigma)$). More precisely, if M is a plausibility-access model with access set function A , then the operation of σ -inference on knowledge produces the model $M_{\hookrightarrow_{\sigma}^K}$, differing from M just in the access set function A' , which is given by

$$A'(w) := \begin{cases} A(w) \cup \{\text{cn}(\sigma)\} & \text{if } (M, w) \Vdash K_{\text{Ex}}\sigma \wedge K_{\text{Ex}}\text{pm}(\sigma) \\ A(w) & \text{otherwise} \end{cases}$$

A new modality $\langle \hookrightarrow_{\sigma}^K \rangle$ is introduced to express the effects of this operation, and its semantic definition is given by

$$(M, w) \Vdash \langle \hookrightarrow_{\sigma}^K \rangle \varphi \quad \text{iff} \quad (M, w) \Vdash K_{\text{Ex}}\sigma \wedge K_{\text{Ex}}\text{pm}(\sigma) \quad \text{and} \quad (M_{\hookrightarrow_{\sigma}^K}, w) \Vdash \varphi$$

But take a closer look at the inference on knowledge operation. What it actually does is to discard all worlds where $K_{\text{Ex}}\sigma \wedge K_{\text{Ex}}\text{pm}(\sigma)$ holds, and replace them with copies that are almost identical, the only difference being their A -sets that, after the operation, will have $\text{cn}(\sigma)$. And this is reasonable because, under the assumption that knowledge is true information, inference based on a known (therefore truth-preserving) rule with known (therefore true) premises is simply deductive reasoning: the premises are true and the rule preserves the truth, so the conclusion *should* be true. In fact, inference based on a known rule with known premises is the act of recognizing two things. First, since the applied rule is truth-preserving and its premises are true, its conclusion *must* be true; and second, situations where the premises are true but the conclusion is not *are not possible*.

The case of beliefs is different, as suggested in [6]. An inference on beliefs is based on a rule that is believed to be truth-preserving, but that it is not necessarily so. Even though it is reasonable to consider a situation in which the premises and the conclusion hold, the agent should not discard a situation where the premises hold but the conclusion does not.

Our proposal is the following. An inference on beliefs should create two copies of each world where the rule and the premises are believed: an exact copy of the original one, and another extending it by adding the rule's conclusion to it. But not only that. The agent believes that the rule is truth-preserving and the premises are true, so the extended world should be more plausible than the 'conclusionless' one.

But, how to create copies of a possible world? We can use the *action models* and *product update* of the so-called *BMS* approach [2].

4.1 Plausibility-access action models

The main idea behind *action models* [2] is that actions can be represented with a model similar to that used for representing the static situation. In other words, just as the agent can be uncertain about which one is the real world, she can also be uncertain about which action has taken place. Then, the uncertainty of the agent *after an action* is a combination of her uncertainty about the situation *before the action* and her uncertainty *about the action* itself.

This idea has been extended in two different directions: in order to deal with plausibility models [3] and in order to deal with non-omniscient multi-agent situations [10]. Our proposal combines and extends these two ideas, now with the aim of dealing with single-agent inference on beliefs. We start by defining the structures that will represent this kind of actions.

Definition 4.1 (Plausibility-access action model). A *plausibility-access action model* is a tuple $A = \langle S, \preceq, \text{Pre}, \text{Pos}_A, \text{Pos}_R \rangle$ where

- $\langle S, \preceq, \text{Pre} \rangle$ is a plausibility action model [3] with S a finite non-empty set of *events*, \preceq a *plausibility* relation on S (with the same requirements as those for a plausibility-access model) and $\text{Pre} : S \rightarrow \mathcal{L}_f$ a *precondition* function indicating the requirement for each event to be executed.
- $\text{Pos}_A : (S \times \wp(\mathcal{L}_f)) \rightarrow \wp(\mathcal{L}_f)$ is the *new access set* function, which will allow us to define the *access set* of the agent in the model that will result from applying this action.
- $\text{Pos}_R : (S \times \wp(\mathcal{L}_r)) \rightarrow \wp(\mathcal{L}_r)$ is the *new rule set* function, which will allow us to define the *rule set* of the agent in the model that will result from applying this action.

Just as before, the plausibility relation \preceq defines an equivalence relation by putting it together with its converse: $\approx := \preceq \cup \succeq$. A pointed plausibility-access action model (A, s) has a distinguished event $s \in S$.

Examples of plausibility-access action models will be shown in Section 4.2. But first we will define the plausibility-access model that results from an action model application as well as the formula that will represent this operation and its semantic interpretation.

Definition 4.2 (Product update). Let $M = \langle W, \leq, V, A, R \rangle$ be a PA model and $A = \langle S, \leq, \text{Pre}, \text{Pos}_A, \text{Pos}_R \rangle$ be a PA action model. The PA model $M \otimes A = \langle W', \leq', V', A', R' \rangle$ is given by

- $W' := \{(w, s) \in (W \times S) \mid (M, w) \models \text{Pre}(s)\}$
- $(w_1, s_1) \leq' (w_2, s_2)$ iff $(s_1 < s_2 \text{ and } w_1 \approx w_2)$ or $(s_1 \approx s_2 \text{ and } w_1 \leq w_2)$
- $V'(w, s) := V(w)$
- $A'(w, s) := \text{Pos}_A(s, A(w))$
- $A'(w, s) := \text{Pos}_R(s, R(w))$

Note how the set of worlds of the new plausibility-access model is given by the restricted cartesian product of W and S ; a pair (w, s) will be a world in the new model iff event s can be executed at world w . The new plausibility order follows the so-called ‘Action-priority’ rule [3], making (w_2, s_2) more plausible than (w_1, s_1) iff either s_2 is strictly more plausible than s_1 and w_1, w_2 are indistinguishable, or else s_1, s_2 are indistinguishable and w_2 is more plausible than w_1 .

Now, for the valuations of the new worlds. First, a new world inherits the atomic valuation of its static component, that is, an atom p holds at (w, s) iff p holds at w . The cases for access sets gives us full generality: the access set of world (w, s) is given by the function Pos_R with the event s and the access set of w as parameters [10]. The case for rule sets is similar.

It is not hard to verify that the *product update* operation preserves plausibility-access models.

Proposition 2. *If M is a plausibility-access model and A a plausibility-access action model, then $M \otimes A$ is a plausibility-access model.*

In order to express how product updates affect the agents’ information, we extend our language with modalities for each pointed plausibility-access action model (A, s) , allowing us to build formulas of the form $\langle A, s \rangle \varphi$, whose semantic interpretation is given below.

Definition 4.3. Let (M, w) be a pointed PA model and let (A, s) be a pointed PA action model with Pre its precondition function.

$$(M, w) \models \langle A, s \rangle \varphi \quad \text{iff} \quad (M, w) \models \text{Pre}(s) \text{ and } (M \otimes A, (w, s)) \models \varphi$$

4.2 Plausibility-access action models for basic inference

The action of inference on knowledge can be represented with plausibility-access action models.

Definition 4.4 (Inference on knowledge). Let σ be a rule. The action of *inference on knowledge* is given by the pointed PA action model $(A_{\rightarrow_{\sigma}^k}, s)$ whose definition (left) and relevant diagram (right) are given by

$$\begin{array}{l} \bullet S := \{s\} \quad \bullet \leq := \{(s, s)\} \quad \bullet \text{Pre}(s) := K_{\text{Ex}}\sigma \wedge K_{\text{Ex}}\text{pm}(\sigma) \\ \bullet \text{Pos}_A(s, X) := X \cup \{\text{cn}(\sigma)\} \quad \bullet \text{Pos}_R(s, Y) := Y \end{array} \quad \begin{array}{c} \hline \Downarrow \\ s \quad X \cup \{\text{cn}(\sigma)\} \\ \hline \end{array}$$

But now we can represent more. Following our previous discussion, here is the action model for basic inference on beliefs.

Definition 4.5 (Basic inference on beliefs). Let σ be a rule. The action of *basic inference on belief* is given by the pointed PA action model $(A_{\rightarrow_{\sigma}^B}, s_1)$ whose definition is

$$\begin{array}{l} \bullet S := \{s_1, s_2\} \\ \bullet \leq := \{(s_1, s_1), (s_1, s_2), (s_2, s_2)\} \\ \bullet \begin{cases} \text{Pre}(s_1) := \text{Pre}_{B\sigma} \\ \text{Pre}(s_2) := \text{Pre}_{B\sigma} \end{cases} \end{array} \quad \bullet \begin{cases} \text{Pos}_A(s_1, X) := X \\ \text{Pos}_A(s_2, X) := X \cup \{\text{cn}(\sigma)\} \\ \text{Pos}_R(s_1, Y) := Y \\ \text{Pos}_R(s_2, Y) := Y \end{cases}$$

The precondition is that the agent believes explicitly the rule and its premises, that is,

$$\text{Pre}_{B\sigma} := B_{\text{Ex}}\sigma \wedge B_{\text{Ex}}\text{pm}(\sigma)$$



$$\begin{array}{c} \hline \Leftarrow s_2 \quad X \cup \{\text{cn}(\sigma)\} \\ \uparrow \\ \Leftarrow s_1 \quad X \\ \hline \end{array}$$

The relevant diagram appears on the right.

4.3 Extended inference: an exploration

Plausibility-access action models allow us to represent more than what we have discussed. As observed in [3], a plausibility relation generates a Grove's system-of-spheres, that is, several layers of possible events ordered according to their plausibility. The presented action models for basic inference on beliefs are just those models with two layers, each one of them having just one event, and with the most plausible one being the extended one. But we do not have to restrict ourselves to such kind of inferences.

Action models with more than two layers allow us to represent inference based on rules with more than one conclusion. The action model on the right has three layers, each one containing one event. Event s_1 preserves access sets, s_2 extends them with the first conclusion and s_3 extends them with both conclusions.



And we can do more by using layers with more than one world, like the action model on the left that allows the agent to have $cn_2(\sigma)$ without having $cn_1(\sigma)$.

So far our examples have one characteristic in common. The *new access set* function is monotone, reflecting the optimism of the agent with respect to the conclusion: events that extend **A**-sets are always more plausible.

Definition 4.6 (Plausibility-access action models for optimistic inference). Plausibility-access action models in which, for every event s_1, s_2 ,

$$s_1 \preceq s_2 \quad \text{implies} \quad \text{Pos}_A(s_1, X) \subseteq \text{Pos}_A(s_2, X)$$

are called action models for *optimistic inference*.

But then we can also consider the opposite case. Models with an *anti-monotone new access set* function reflect the pessimism of the agent with respect to the conclusion: events that extend **A**-sets are always less plausible.

Definition 4.7 (Plausibility-access action models for pessimistic inference). Plausibility-access action models in which, for every event s_1, s_2 ,

$$s_1 \preceq s_2 \quad \text{implies} \quad \text{Pos}_A(s_1, X) \supseteq \text{Pos}_A(s_2, X)$$

are called action models for *pessimistic inference*.

Of course these two classes do not cover all possibilities. Plausibility-access action models allow us to represent many different and complex inferences whose detailed study has to be left for further work.

4.4 Brief discussion on completeness

The reduction axioms of [3] are inherited by our system. In particular, the following one states the way the plausibility relation changes:

$$\langle A, s \rangle \langle \leq \rangle \varphi \leftrightarrow \left(\text{Pre}(s) \wedge \left(\bigvee_{s \leq s'} \langle \approx \rangle \langle A, s' \rangle \varphi \vee \bigvee_{s \approx s''} \langle \leq \rangle \langle A, s'' \rangle \varphi \right) \right)$$

But when looking for reduction axioms for access and rule set formulas, Pos_A and Pos_R pose a problem. The reason is that they allow the new access and rule sets to be arbitrary sets. Compare this with other product update definitions. The one of [7] can change the atomic valuation, but the set of worlds in which a given atomic proposition will be true should be given by a formula of the language; the one of [16] can change the relation in a point-wise way, but the new relation is given in terms of the previous ones by using only regular operations. Our current efforts focus on particular definitions expressive enough to describe our desired inferences and restricted enough to get the needed reduction axioms.

5 Conclusions and further work

We have presented a framework for representing implicit and explicit beliefs. We have also provided representations of three actions that modify them, starting with those of *explicit upgrade* and *retraction* but, more important, discussing intuitive ideas and proposing a rich framework for representing the action of *inference on beliefs*.

There are parts of this work that deserve further exploration, the most appealing being the study of the different kind of inferences that we can represent with plausibility-access action models. We have defined those for inference on knowledge and basic inference on beliefs, and we have briefly explored some others, but our structures can represent much more. Another interesting extension is to look at dynamics of rules, that is, to look for reasonable actions that extend not only the rules the agent knows [28], but also the rules she believes. These two studies will not be complete without the appropriate axiom system for our product update definition. Finally we mention a third direction: the study of a multi-agent setting, including not only the addition of more agents to the picture, but also the analysis of implicit/explicit versions of multi-agent notions, like common knowledge and common beliefs.

References

- [1] T. Ågotnes and N. Alechina, editors. *Special issue on Logics for Resource Bounded Agents*, 2009. *Journal of Logic, Language and Information*, 18(1).

- [2] A. Baltag, L. Moss, and S. Solecki. The logic of public announcements, common knowledge and private suspicious. SEN-R9922, CWI, Amsterdam, 1999.
- [3] A. Baltag and S. Smets. A qualitative theory of dynamic interactive belief revision. In G. Bonanno, W. van der Hoek, and M. Wooldridge, editors, *Logic and the Foundations of Game and Decision Theory (LOFT7)*, volume 3 of *Texts in Logic and Games*, pages 13–60. AUP, 2008.
- [4] J. van Benthem. Dynamic logic for belief revision. *Journal of Applied Non-Classical Logics*, 17(2):129–155, 2007.
- [5] J. van Benthem. Merging observation and access in dynamic logic. *Journal of Logic Studies*, 1(1):1–17, 2008.
- [6] J. van Benthem. Logic, mathematics, and general agency. In P. Bour, M. Rebuschi, and L. Rollet, editors, *Festschrift for Gerhard Heinzmann*. Laboratoire d’histoire des sciences et de la philosophie, Nancy, 2009.
- [7] J. van Benthem, J. van Eijck, and B. Kooi. Logics of communication and change. *Information and Computation*, 204(11):1620–1662, 2006.
- [8] J. van Benthem and B. Kooi. Reduction axioms for epistemic actions. In R. Schmidt, I. Pratt-Hartmann, M. Reynolds, and H. Wansing, editors, *Advances in Modal Logic (Technical Report UMCS-04-09-01)*, pages 197–211. University of Manchester, 2004.
- [9] J. van Benthem and F. Liu. Dynamic logic of preference upgrade. *Journal of Applied Non-Classical Logics*, 17(2):157–182, 2007.
- [10] J. van Benthem and F. R. Velázquez-Quesada. Inference, promotion, and the dynamics of awareness. PP-2009-43, ILLC, Universiteit van Amsterdam, 2009.
- [11] P. Blackburn, M. de Rijke, and I. Venema. *Modal Logic*. Cambridge University Press, 2001.
- [12] O. Board. Dynamic interactive epistemology. *Games and Economic Behavior*, 49(1):49–80, 2004.
- [13] H. van Ditmarsch, A. Herzig, J. Lang, and P. Marquis. Introspective forgetting. *Synthese (KRA)*, 169(2):405–423, 2009.
- [14] H. van Ditmarsch, W. van der Hoek, and B. Kooi. *Dynamic Epistemic Logic*, volume 337 of *Synthese Library Series*. Springer, 2007.
- [15] H. N. Duc. *Resource-Bounded Reasoning about Knowledge*. PhD thesis, Institut für Informatik, Universität Leipzig, Leipzig, Germany, 2001.
- [16] J. van Eijck and Y. Wang. Propositional dynamic logic as a logic of belief revision. In W. Hodges and R. J. G. B. de Queiroz, editors, *WoLLIC*, volume 5110 of *LNCS*, pages 136–148. Springer, 2008.
- [17] R. Fagin and J. Y. Halpern. Belief, awareness, and limited reasoning. *Artificial Intelligence*, 34(1):39–76, 1988.

- [18] D. Grossi and F. R. Velázquez-Quesada. *Twelve Angry Men: A study on the fine-grain of announcements*. In X. He, J. F. Horty, and E. Pacuit, editors, *LORI*, volume 5834 of *LNCS*, pages 147–160. Springer, 2009.
- [19] A. Grove. Two modellings for theory change. *Journal of Philosophical Logic*, 17(2):157–170, 1988.
- [20] J. Y. Halpern, editor. *Proceedings of the 1st Conference on Theoretical Aspects of Reasoning about Knowledge, Monterey, CA, March 1986*, San Francisco, CA, USA, 1986. Morgan Kaufmann Publishers Inc.
- [21] J. Hintikka. *Knowledge and Belief: An Introduction to the Logic of the Two Notions*. Cornell University Press, Ithaca, N.Y., 1962.
- [22] M. Jago. Rule-based and resource-bounded: A new look at epistemic logic. In T. Ågotnes and N. Alechina, editors, *Proceedings of the Workshop on Logics for Resource-Bounded Agents*, pages 63–77, Malaga, Spain, 2006.
- [23] K. Konolige. Belief and incompleteness. Technical Report 319, SRI, 1984.
- [24] G. Lakemeyer. Steps towards a first-order logic of explicit and implicit belief. In Halpern [20], pages 325–340.
- [25] H. J. Levesque. A logic of implicit and explicit belief. In *Proc. of AAAI-84*, pages 198–202, Austin, TX, 1984.
- [26] K. Segerberg. The basic dynamic doxastic logic of AGM. In M.-A. Williams and H. Rott, editors, *Frontiers in Belief Revision*, number 22 in Applied Logic Series, pages 57–84. Kluwer Academic Publishers, 2001.
- [27] M. Y. Vardi. On epistemic logic and logical omniscience. In Halpern [20], pages 293–305.
- [28] F. R. Velázquez-Quesada. Inference and update. *Synthese (KRA)*, 169(2):283–300, 2009.