



HAL
open science

Raising awareness without disclosing truth

Line van den Berg, Manuel Atencia, Jérôme Euzenat

► **To cite this version:**

Line van den Berg, Manuel Atencia, Jérôme Euzenat. Raising awareness without disclosing truth. *Annals of Mathematics and Artificial Intelligence*, 2023, 91 (4), pp.431-464. 10.1007/s10472-022-09809-y . hal-04351134

HAL Id: hal-04351134

<https://hal.science/hal-04351134>

Submitted on 18 Dec 2023

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution 4.0 International License

Raising Awareness without Disclosing Truth

Line van den Berg · Manuel Atencia ·
Jérôme Euzenat

Received: 2021-02-14 / Accepted: 2022-07-22 / Published: 2022-10-27

Abstract Agents use their own vocabularies to reason and talk about the world. Public signature awareness is satisfied if agents are aware of the vocabularies, or signatures, used by all agents they may, eventually, interact with. Multi-agent modal logics and in particular Dynamic Epistemic Logic rely on public signature awareness for modeling information flow in multi-agent systems. However, this assumption is not desirable for *dynamic* and *open* multi-agent systems because (1) it prevents agents to use unique signatures other agents are unaware of, (2) it prevents agents to openly extend their signatures when encountering new information, and (3) it requires that all future knowledge and beliefs of agents are bounded by the current state.

We propose a new semantics for awareness that enables us to drop public signature awareness. This semantics is based on partial valuation functions and weakly reflexive relations. Dynamics for raising public and private awareness are then defined in such a way as to differentiate between becoming aware of a proposition and learning its truth value. With this, we show that knowledge and beliefs are not affected through the raising operations.

Keywords Awareness · Raising awareness · Dynamic epistemic logic · Partial valuations · Multi-agent systems

L. van den Berg
Univ. Grenoble Alpes, Inria, CNRS, Grenoble INP, LIG, F-38000 Grenoble France. E-mail:
line.van-den-berg@inria.fr

M. Atencia
Univ. Grenoble Alpes, Inria, CNRS, Grenoble INP, LIG, F-38000 Grenoble France. E-mail:
manuel.atencia@inria.fr

J. Euzenat
Univ. Grenoble Alpes, Inria, CNRS, Grenoble INP, LIG, F-38000 Grenoble France. E-mail:
jerome.euzenat@inria.fr

1 Introduction

Agents use propositions to represent the information they have about the world, that together form their vocabulary, also called their *signature*. Agents' signatures may be the same or different, and in the latter case either exclusive or overlapping. When they are the same and this is known by all agents, *public signature awareness* holds: it is commonly known that agents use the same propositions—i.e. they are *aware* of the same vocabulary. Multi-agent modal logics, and in particular Dynamic Epistemic Logic (DEL), are frameworks for modeling the information flow in multi-agent systems and are based on this assumption. DEL has been used as a framework to investigate how knowledge and beliefs of agents evolve under dynamic upgrades in multi-agents systems and has been applied to communication [8,19], belief revision [6] and agent interaction [7]. However, because of public signature awareness, in DEL agents have full awareness of *all* the propositions used currently, *and* in the future, by *all* agents. This ensures that any formula that is communicated, via the dynamic upgrades, can be understood by all agents and that each agent can update her knowledge and beliefs accordingly. But as a consequence, all future evolutions of agents' knowledge and beliefs are constrained to the current situation: learning takes place via eliminating or re-organizing the possibilities agents consider, rather than extending it with the new information.

The public signature awareness assumption becomes a restriction when considering *dynamic* and *open* multi-agent systems. In such systems, agents may use different signatures to communicate and may be required to continuously adapt and extend their signatures to their environment [2, 20, 21]. A typical question for such systems is: how do agents understand each other and interact if not everyone is aware of the propositions used?

The problem of public signature awareness is hidden in the semantics of DEL: valuation functions are total functions and accessibility relations are reflexive. Therefore, in order to eliminate public signature awareness, we need to adjust the structure of the models. In this paper, we propose a novel semantics, called Partial Dynamic Epistemic Logic (ParDEL), with valuations that are *partial* (instead of total) and relations that are *weakly reflexive* (instead of reflexive). The first adjustment creates a distinction between *uncertain* agents, agents that are aware of a proposition but do not know its truth value, and *unaware* agents, agents that do not consider the proposition in its entirety. The second adjustment enables agents to use different signatures to represent their knowledge and beliefs.

We extend this semantics with raising awareness dynamics that allow agents extend their signatures when encountering new information from the environment or from other agents they interact with. These operations act on the models by duplicating the worlds in which the proposition, whose awareness is raised of, is undefined, making it true in one while false in the other. This ensures that raising awareness does not imply disclosing truth. We then show that raising awareness preserves agents' knowledge and beliefs and that, if a sentence they know or believe after raising awareness of a proposition does not involve this proposition, this sentence was already known or believed before the operation. This result is extended to non-epistemic formulas and it is proven that raising awareness of a proposition does not imply disclosing its truth value. Hence, awareness and truth are disconnected.

The semantics introduced in this paper can be used to model communication and interaction between agents that use different signatures and extend their signatures when encountering new propositions via raising awareness. As such, it is a generalization of DEL for modeling dynamic and open multi-agent systems.

In the remainder, we discuss the related work (§2) and introduce Dynamic Epistemic Logic (§3). We then discuss how we characterize awareness and unawareness of agents through partial valuations (§4). We formally define the syntax and semantics of Partial Dynamic Epistemic Logic (§5). Then, we formalize the dynamics for raising public and private awareness on ParDEL+ (§6) and show that knowledge and beliefs of agents are not affected through the raising operations (§7). Finally, we discuss issues about supporting a dual forgetting awareness operator (§8) and we conclude by pointing out the applications of this work and discuss possible future directions (§9).

2 Related work

Dynamic Epistemic Logics (DEL) are a family of modal logics (see e.g. [12]) studying knowledge, belief and other epistemic attitudes that are studied in epistemic logic [23], under model change, using formal languages and mathematical models [19]. These model changes, also called dynamic upgrades, are model-transforming actions that enable to shift away from the static semantics of truth in an individual Kripke model to a dynamic semantics of truth occurring across Kripke models, specified by the upgrade. This allows us to analyze the epistemic and doxastic consequences of, for example, public announcements [35]. DEL has been widely used as a formal framework to model agent communication [8, 19, 35], belief revision [6] and agent interaction [7].

Awareness has been first formalized in logic by Fagin and Halpern [22] as a way to solve the problem of logical omniscience [30]. They wondered how agents can say that they know or believe a proposition p if p is a concept they are completely unaware of. As a consequence, the interpretation of $K_a\phi$ is changed from “agent a knows ϕ ” to “agent a *implicitly* knows ϕ ”, denoting the knowledge an agent *could* eventually get. Another notion of knowledge called *explicit knowledge* is then defined as a combination of implicit knowledge and awareness, denoting the knowledge an agent has access to at a particular moment. The logic introduced in [22], called Awareness Logic, extends the language of epistemic logic with an operator $A_a\phi$ that reads as “agent a is aware of ϕ ” and that is interpreted with respect to an awareness function assigning to each world and each agent a set of sentences. Agents then (explicitly) know or believe a formula ϕ if it is true in all accessible worlds or all most plausible worlds, respectively, such that $A_a\phi$ holds. The awareness function can therefore be viewed as acting like a filter: extracting explicit knowledge from implicit knowledge.

With this notion of awareness, an interest has arisen to study the dynamics of awareness using DEL, e.g. [9, 14–16, 26, 28, 29, 34]. Particularly interesting to this paper, in [14, 15] Awareness Logic is extended to account for dynamic awareness, where the dynamic part is modeled by a bisimulation quantification on structures. Event models for changing awareness following this approach are defined in [16, 17] and modalities to change awareness have been introduced in [9, 15], called the *consider* and *drop* operations, that, respectively, extend or reduce the scope of the

awareness function. A complete dynamic epistemic logic of awareness is defined and the operations are generalized to multi-agent situations where awareness changes may occur privately in [9, 17].

These works on awareness have mainly concentrated on awareness of the truth value of a statement not on awareness of the statement itself. This is because they are, like DEL, based on total valuation functions and only awareness is considered as a partial function. This means that raising awareness comes equipped with disclosing the ‘underlying’ truth value of the proposition awareness is raised of [14, 15]. The underlying truth value may be defined as such that agents become ignorant when raising their awareness [17], however, the problem is that this needs to be determined in advance. That is, the valuations of the proposition awareness is raised of are already defined, but are ‘invisible’ to the agents. Therefore, two problems remain: it disables agents to openly evolve their signatures when encountering new information and any future evolution of agents’ knowledge and beliefs, and now also awareness, is bound by the initial setting.

In the notion of awareness introduced in this paper, awareness is present in the structure of the logic and no awareness operator $A\phi$ or awareness function is required. Instead, awareness arises from the use of *partial valuation functions* and *weakly reflexive relations*. This means that agents may use different signatures, but we also tackle the other two problems: (1) their signatures can openly evolve via raising awareness causing knowledge, belief and awareness to evolve without any prior defined way of how this evolution might take place, and (2) awareness is completely disconnected from truth: raising awareness does not imply disclosing its truth values.

The connection between partial valuation functions and awareness is novel in this paper. Partial valuation functions have been introduced for (Dynamic) Epistemic Logic in [31, 32, 41], where worlds of the models are equipped with partial, instead of total, valuation functions that either interpret the propositions as true or false or do not interpret them, leaving them *undefined*. This offers a more natural way to deal with growth of information by extending the models rather than merely reducing or re-organizing them, what is the typical approach in standard Dynamic Epistemic Logic to model changes in information [19]. This means that in DEL only certainty grows in parallel with new information, but in a partial approach also awareness can grow in parallel with new information. Yet, a link to awareness, and therefore the possibility of agents using different, dynamic signatures to represent their knowledge and beliefs, has not yet been established.

We extend the work by [27, 31, 32, 41] in three ways: (1) we replace reflexivity by *weak reflexivity*, (2) we consider a different clause for the falsification of conjunctions and (3) we make a connection between partial valuations and awareness of agents. Weak reflexivity enables us to model how agents with different signatures, other agents are unaware of, interact. The new clause for falsification of conjunctions enforces that both conjuncts must belong to the domain of the valuation functions at the world considered in order to be false. This ensures that agents can only know that a conjunction “ p and q ” is false, if they are aware of both p and q and know that at least one of the two conjuncts is false. Finally, the connection with awareness, gives a novel semantics for awareness and unawareness of agents in which becoming aware of a proposition and learning its truth value are two independent acts and in which agents can extend their signatures when learning new information from the environment or from other agents.

3 Dynamic Epistemic Logic

We consider the syntax and semantics of Dynamic Epistemic Logic as introduced by [4].

Definition 1 (Syntax of DEL) Given a countable, non-empty set P of propositions and a finite, non-empty set \mathcal{A} of agents, the *syntax*, \mathcal{L}_{DEL} , of (multi-agent) Dynamic Epistemic Logic is defined in the following way:

$$\phi ::= p \mid \phi \wedge \psi \mid \neg\phi \mid K_a\phi \mid B_a\phi \mid [\dagger\phi]\psi$$

where $p \in P$ is a proposition, K_a and B_a are the knowledge and belief operators for each agent $a \in \mathcal{A}$ and $\dagger\phi$ with $\dagger \in \{!, \uparrow, \downarrow\}$ the dynamic upgrades.

The connectives \vee and \rightarrow , and the duals $\hat{K}_a, \hat{B}_a, \langle \dagger\phi \rangle$ are defined in the usual way: $\phi \vee \psi$ iff $\neg(\neg\phi \wedge \neg\psi)$, $\phi \rightarrow \psi$ iff $\neg\phi \vee \psi$, $\hat{K}_a\phi$ iff $\neg K_a\neg\phi$, $\hat{B}_a\phi$ iff $\neg B_a\neg\phi$, and $\langle \dagger\phi \rangle\psi$ iff $\neg[\dagger\phi]\neg\psi$.

We read the formula $K_a\phi$ as “agent a knows that ϕ is true” and the formula $B_a\phi$ as “agent a believes that ϕ is true”. The standard semantics for \mathcal{L}_{DEL} are given by means of Kripke, or DEL, models with plausibility relations.

Definition 2 (DEL Model) Given a countable, non-empty set P of propositions and a finite, non-empty set \mathcal{A} of agents, a *model* of (multi-agent) Dynamic Epistemic Logic is a triple $\mathcal{M} = \langle W, (\leq_a)_{a \in \mathcal{A}}, V \rangle$ where

- W is a non-empty set of *worlds*;
- $(\leq_a)_{a \in \mathcal{A}} : \mathcal{A} \rightarrow \mathcal{P}(W \times W)$ are the plausibility relations on W , one for each agent, that are well-founded, locally connected preorders;
- $V : P \rightarrow \mathcal{P}(W)$ is a propositional valuation mapping propositions to sets of worlds in which that proposition is true.

A *pointed DEL model* is a pair $\langle \mathcal{M}, w \rangle$, with $w \in W$.

For a DEL model $\mathcal{M} = \langle W, (\leq_a)_{a \in \mathcal{A}}, V \rangle$, the part $\langle W, (\leq_a)_{a \in \mathcal{A}} \rangle$ is also referred to as a *DEL frame* and is denoted by \mathcal{F} . By convention, $W^{\mathcal{M}}$, $R_a^{\mathcal{M}}$ and $V^{\mathcal{M}}$ are used to refer to the components of \mathcal{M} , but we also omit the superscript \mathcal{M} if it is clear from the context which model we are concerned with. We write $w \in \mathcal{M}$ to mean $w \in W^{\mathcal{M}}$. Furthermore, we write $V_w(p) = 1$ to denote that $w \in V(p)$, and $V_w(p) = 0$ to denote that $w \notin V(p)$.

The plausibility relation $w \leq_a v$ reads as “ v is at least as plausible as w for agent a ”. We write $w \geq_a v$ if and only if $v \leq_a w$. This plausibility relation is used to define knowledge and beliefs of agents: agent a knows something if and only if it is true at all the worlds agent a considers plausible (in any direction) and agent a believes something if and only if it is true at the maximal worlds with respect to the plausibility relation for agent a . Knowledge and beliefs can also be defined with respect to the epistemic (\sim_a) and doxastic relations (\rightarrow_a) that can be deduced from the plausibility relation \leq_a .

Definition 3 (Epistemic and Doxastic Relation) Given a DEL model $\mathcal{M} = \langle W, (\leq_a)_{a \in \mathcal{A}}, V \rangle$ for a set \mathcal{A} of agents, then the epistemic relation \sim_a is defined as:

$$w \sim_a v \text{ iff } w (\leq_a \cup \geq_a)^* v \tag{1}$$

And the doxastic relation \rightarrow_a is defined as:

$$w \rightarrow_a v \text{ iff } v \in \text{Max}_{\leq_a} |w|_a \quad (2)$$

where R^* is the transitive closure of any relation R and $|w|_a$ is the *information cell* (or *accessible cell*) of agent a at state w and is defined by:

$$|w|_a = \{v \in W \mid w \sim_a v\} \quad (3)$$

Then an agent a knows ϕ if it holds at all the worlds reached via \sim_a and she believes ϕ if it holds at all the worlds reached via \rightarrow_a .

It follows from the properties of \leq_a and \geq_a that the relations \sim_a are reflexive, transitive and symmetric, and the relations \rightarrow_a are transitive, serial and Euclidean. Therefore they satisfy the usual properties of knowledge and beliefs, *S5* (K (distributivity), T (factivity), 4 (positive introspection) and 5 (negative introspection)) and *KD45* (K , D (consistency), 4 and 5), respectively [19].

The dynamic upgrades $!\phi$, $\uparrow\phi$ and $\dagger\phi$ act as model transformers. Model transformers are functions whose domain and range are the set of DEL models.

Definition 4 (Model Transformer) A *model transformer* is a function $\dagger\phi : \mathcal{M} \mapsto \mathcal{M}^{\dagger\phi}$, applying a certain action to \mathcal{M} to obtain $\mathcal{M}^{\dagger\phi} = \langle W^{\dagger\phi}, (\leq_a^{\dagger\phi})_{a \in \mathcal{A}}, V^{\dagger\phi} \rangle$. We consider three model transformers $!\phi$, $\uparrow\phi$ and $\dagger\phi$ that are defined as follows, with $\|\phi\|_{\mathcal{M}} := \{w \in W \mid \mathcal{M}, w \models \phi\}$ denoting the set of worlds in which ϕ is true:

Announcement ($!\phi$) deletes all ' $\neg\phi$ '-worlds from the model, i.e. $W^{!\phi} = \|\phi\|_{\mathcal{M}}$, $w \leq_a^{!\phi} v$ iff $w \leq_a v$ and $w, v \in W^{!\phi}$, $V^{!\phi}(p) = V(p) \cap \|\phi\|_{\mathcal{M}}$;

Radical upgrade ($\uparrow\phi$) makes all ϕ worlds more plausible than all $\neg\phi$ worlds, and within these two zones, the old ordering remains. I.e. $W^{\uparrow\phi} = W$, $w \leq_a^{\uparrow\phi} v$ iff $v \in \|\phi\|_{\mathcal{M}}$ and $w \in \|\neg\phi\|_{\mathcal{M}}$ or else if $w \leq_a v$, and $V^{\uparrow\phi}(p) = V(p)$;

Conservative upgrade ($\dagger\phi$) makes the best ' ϕ '-worlds more plausible than all other worlds, while the old ordering on the rest of the worlds remains. I.e. $W^{\dagger\phi} = W$, $w \leq_a^{\dagger\phi} v$ iff either $v \in \text{Max}_{\leq_a}(|w|_a \cap \|\phi\|_{\mathcal{M}})$ or $w \leq_a v$, $V^{\dagger\phi}(p) = V(p)$.

Satisfiability is considered with respect to a pointed model $\langle \mathcal{M}, w \rangle$.

Definition 5 (Satisfiability for DEL) *Satisfiability* for Dynamic Epistemic Logic by a pointed model $\langle \mathcal{M}, w \rangle$ is defined in the following way:

$$\begin{array}{ll} \mathcal{M}, w \models p & \text{iff } w \in V(p) \\ \mathcal{M}, w \models \phi \wedge \psi & \text{iff } \mathcal{M}, w \models \phi \text{ and } \mathcal{M}, w \models \psi \\ \mathcal{M}, w \models \neg\phi & \text{iff } \mathcal{M}, w \not\models \phi \\ \mathcal{M}, w \models K_a\phi & \text{iff } \forall v \text{ s.t. } w \sim_a v : \mathcal{M}, v \models \phi \\ \mathcal{M}, w \models B_a\phi & \text{iff } \forall v \text{ s.t. } w \rightarrow_a v : \mathcal{M}, v \models \phi \\ \mathcal{M}, w \models [\dagger\phi]\psi & \text{iff } \mathcal{M}^{\dagger\phi}, w \models \psi \end{array}$$


where $\dagger \in \{!, \uparrow, \dagger\}$.

We also write $\dagger_1\phi; \dagger_2\psi$ for the sequence of upgrades $\dagger_1\phi$ and then $\dagger_2\psi$ meaning that first $\dagger_1\phi$ is applied and then $\dagger_2\psi$, i.e. the model $\mathcal{M}^{\dagger_1\phi; \dagger_2\psi}$ is defined as $(\mathcal{M}^{\dagger_1\phi})^{\dagger_2\psi}$.

We may use \perp to denote $p \wedge \neg p$ for any proposition $p \in P$, which is false on all (non-empty) models, and likewise \top to denote $\neg \perp$, which is true on all (non-empty) models. As usual, we say that a set of formulas is *consistent* if there is a pointed model satisfying all formulas of the set. Otherwise, a set of formulas is *inconsistent*. Furthermore, a formula ϕ is a *consequence* of a set of formulas Γ (written $\Gamma \models \phi$) if every pointed model $\langle \mathcal{M}, w \rangle$ satisfying all formulas of Γ , also satisfy ϕ .


When we draw DEL models, we write, inside the worlds, p to denote that the valuation of p at that world is 1 and \bar{p} to denote that the valuation is 0. Relations $w \leq_a v$ are represented by drawing an arrow from w to v ($w \rightarrow v$) with the label a . When considering pointed models $\langle \mathcal{M}, w \rangle$ we also double circle the world w . Consider the following example of a DEL model.

Example 1 A DEL model $\mathcal{M} = \langle W, (\leq_a)_{a \in \mathcal{A}}, V \rangle$ for a set $\mathcal{A} = \{a, b\}$ of agents is depicted below. In the pointed model $\langle \mathcal{M}, w \rangle$ agent a knows that p is true but q is false, and agent b believes that q is true.



The different upgrades are used to capture differences in the trustworthiness of the information source: the source may be considered as an infallible source (announcements), as a highly trusted, but fallible source (radical upgrades), or as a trusted, but not infallible source (conservative upgrades) [3]. Radical and conservative upgrades have been discussed under various names in the context of Belief Revision (for example the AGM framework [24]), e.g. in [38, 42], and have been formalized for DEL in an attempt to bridge DEL with Belief Revision in [6]. Fixed points of announcements, radical and conservative upgrades have been investigated in [5].

Example 2 Let $\mathcal{M} = \langle W, (\leq_a)_{a \in \mathcal{A}}, V \rangle$ be the DEL model of Example 1. Then $\mathcal{M}^{!p}$, $\mathcal{M}^{\uparrow p}$ and $\mathcal{M}^{\uparrow p}$ are depicted below. It holds that, among others, $\mathcal{M}^{!p}, w \models K_b p$, $\mathcal{M}^{\uparrow p}, w \models B_b p$ and $\mathcal{M}^{\uparrow p}, w \models B_b p$.



Interestingly, announcements, radical and conservative upgrades do not actually add something new to the language. Indeed, the expressivity of DEL is the same as for an epistemic-doxastic logic: any sentence involving an announcement, radical or conservative upgrade can be reduced to one without via *reduction axioms* [6, 35].

Lastly, when working with DEL models, a useful and important notion is that of *bisimulation*. It formalizes when two models are semantically equivalent.

Definition 6 (Bisimulation) Let two DEL models $\mathcal{M} = \langle W, (\leq_a)_{a \in \mathcal{A}}, V \rangle$ and $\mathcal{M}' = \langle W', (\leq'_a)_{a \in \mathcal{A}}, V' \rangle$ be given for set \mathcal{A} of agents.

A relation $Z \subseteq W \times W'$ is a *bisimulation* if and only if for all $(w, w') \in Z$ the following three conditions hold:

- **[Propositional agreement]** $V(w) = V'(w')$;
- **[Forth]** For every agent $a \in \mathcal{A}$ and for every $v \in W$ such that $w \leq_a v$ there exists a $v' \in W'$ such that $w' \leq'_a v'$ and $(v, v') \in Z$;
- **[Back]** For every agent $a \in \mathcal{A}$ and for every $v' \in W'$ such that $w' \leq'_a v'$ there exists a $v \in W$ such that $w \leq_a v$ and $(v, v') \in Z$.

Two pointed models $\langle \mathcal{M}, w \rangle$ and $\langle \mathcal{M}', w' \rangle$ are *bisimilar* if and only if there is a bisimulation Z such that $(w, w') \in Z$.

In DEL, the semantic notion of bisimulation coincides with the statement that models satisfy the same formulas [19].

3.1 Dynamic Epistemic Logic with Event Models

In the previous section, three ways have been discussed and formalized to change the knowledge and beliefs of agents: announcements, radical upgrades and conservative upgrades. In this section, we look at a generalization of these upgrades called *event models*, or *action models*, as introduced in [3]. Event models are relational structures that allow us to talk about the dynamics of information in the same way that DEL models formalize static information. The general idea is to think of event models as Kripke models, but instead of consisting of worlds, we consider

it to be consisting of a set of *events*, and instead of a valuation a *precondition* is defined. Like on a DEL model, the relational structure of an event model specifies which events the agents can tell apart.

Event models can be used to describe a variety of informational events: from public announcements to more subtle communication containing privacy, misleading or suspicion. For example, information may be shared in secret, hidden completely (other agents do not observe the communication) or partially (other agents observe the communication, but not what is communicated) from others.

To formalize these situations, we first introduce an alternative language: the language of Epistemic Action Logic (EAL) [3].

Definition 7 (Syntax of EAL) Given a countable, non-empty set P of propositions and a finite, non-empty set \mathcal{A} of agents, the *syntax*, \mathcal{L}_{EAL} , of (multi-agent) Epistemic Action Logic (EAL) is defined in the following way.

$$\phi ::= p \mid \phi \wedge \psi \mid \neg\phi \mid K_a\phi \mid B_a\phi \mid [\langle \mathcal{E}, e \rangle]\phi$$

where $p \in P$ is a proposition, K_a and B_a are the knowledge and belief operators for each agent a , and $\langle \mathcal{E}, e \rangle$ are pointed event models.

Models of EAL are equivalent to models of DEL (Definition 2), with a plausibility relation for each agent that is a well-founded, locally connected preorder.

Event models for EAL provide a relational structure to dynamic upgrades. Formally, an event model is like a Kripke model but instead of worlds we consider *events* and instead of a valuation a *precondition* is defined.

Definition 8 (Event Model) Let \mathcal{A} be a finite, non-empty set of agents. An *event model* for EAL is a triple $\mathcal{E} = \langle E, (R_a)_{a \in \mathcal{A}}, pre \rangle$ where

- E is a non-empty, finite set of *events*;
- $(R_a)_{a \in \mathcal{A}} \subseteq E \times E$ are the *accessibility relations* on E , one for each agent $a \in \mathcal{A}$;
- $pre : E \rightarrow \mathcal{L}_{DEL}$ is a *precondition function* assigning to each event a formula ϕ .

Given an event model \mathcal{E} and an event $e \in E$, $\langle \mathcal{E}, e \rangle$ is a *pointed event model*.

When drawing a pointed event model $\langle \mathcal{E}, e \rangle$, events are drawn as squares to distinguish them from EAL models and e is double-squared.

The *product update* $\mathcal{M} \otimes \mathcal{E}$ determines what happens if an event model \mathcal{E} takes place on a EAL model \mathcal{M} [3].

Definition 9 (Product Update) Let $\mathcal{M} = \langle W, (\leq_a)_{a \in \mathcal{A}}, V \rangle$ be a EAL model and $\mathcal{E} = \langle E, (R_a)_{a \in \mathcal{A}}, pre \rangle$ be an event model. Their *product update*, denoted by $\mathcal{M} \otimes \mathcal{E}$, is the triple $\langle W^{\mathcal{M} \otimes \mathcal{E}}, (\leq_a^{\mathcal{M} \otimes \mathcal{E}})_{a \in \mathcal{A}}, V^{\mathcal{M} \otimes \mathcal{E}} \rangle$ defined by:

- $W^{\mathcal{M} \otimes \mathcal{E}} = \{ \langle w, e \rangle \in W \times E \mid \mathcal{M}, w \models pre(e) \}$
- $\langle w, e \rangle \leq_a^{\mathcal{M} \otimes \mathcal{E}} \langle w', e' \rangle$ iff $\langle w, e \rangle, \langle w', e' \rangle \in W^{\mathcal{M} \otimes \mathcal{E}}$, $w \leq_a w'$ and $e R_a e'$
- $V^{\mathcal{M} \otimes \mathcal{E}}(p) = \{ \langle w, e \rangle \in W \times E \mid w \in V(p) \}$

The product update $\mathcal{M} \otimes \mathcal{E}$ is the result of the events $e \in E$ happening at the worlds $w \in W$ whenever w satisfies the precondition $pre(e)$. The precondition therefore serves as a selection to which worlds an event may be applied. For example, if $pre(e) = \phi$, that means that the event e may only be applied to worlds

w that make ϕ true. Then, if e is the sole event of \mathcal{E} , this means that the worlds falsifying ϕ are deleted from the product update. In the following we also refer to the events e such that $\mathcal{M}, w \models pre(e)$ as the events that *can be applied* to w .

The accessibility relations R_a express how the different agents observe the event. This determines which relations remain in the product update from the initial epistemic model. Hence, for a to have access to a world in the product update, there needs to be a \leq_a -arrow between the corresponding worlds in the DEL model and a R_a -arrow between their corresponding events in the event model. In general, agents may observe the event differently.

Finally, the valuation in the product update is the same as before, ranging over the worlds in the product update.

Satisfiability for events is then determined with respect to the product update.

Definition 10 (Satisfiability for Events) *Satisfiability* for EAL extends satisfiability for DEL (Definition 5), replacing the last clause by:

$$\mathcal{M}, w \models [\langle \mathcal{E}, e \rangle] \psi \text{ iff } \mathcal{M}, w \models pre(e) \text{ implies that } \mathcal{M} \otimes \mathcal{E}, \langle w, e \rangle \models \psi$$

We can capture public announcements with event models.


Definition 11 (Public Announcement) The pointed event model for the public announcement $!\phi$ is $\langle \mathcal{E}_{!\phi}, e_\phi \rangle$ where $\mathcal{E}_{!\phi} = \langle \{e_\phi\}, (I_a)_{a \in \mathcal{A}}, pre \rangle$ and $pre(e_\phi) = \phi$:



Indeed, the event model for public announcements is equivalent to the model transformer $!\phi$ in Definition 5: only the worlds satisfying the precondition, namely ϕ , remain in the resulting model, while the accessibility relations to and from these worlds are equivalent to the relations in the initial model. In other words, $\neg\phi$ -worlds are deleted.

Giving event structures to dynamic upgrades is interesting because it allows for more complex upgrades such as private announcements. A private announcement is an announcement that is only received by a subset G of agents and that is “invisible” to the other agents [3].

Definition 12 (Private Announcement) The pointed event model for the fully private announcement $!_G\phi$ to a group $G \subseteq \mathcal{A}$ is $\langle \mathcal{E}_{!_G\phi}, e_\phi \rangle$ where the event model is defined as $\mathcal{E}_{!_G\phi} = \langle \{e_\phi, e_\top\}, (R_a)_{a \in \mathcal{A}}, pre \rangle$ such that for $a \in G$, $R_a = \{ \langle e_\phi, e_\phi \rangle, \langle e_\top, e_\top \rangle \}$ and otherwise $R_a = \{e_\phi, e_\top\} \times \{e_\top\}$, $pre(e_\phi) = \phi$ and $pre(e_\top) = \top$:



A private announcement of a formula ϕ to a group G of agents makes two copies of the worlds: in one of the copies, only accessible by agents in the group G , all the worlds in which ϕ is false are eliminated (through e_ϕ), leaving these agents to learn ϕ , whereas the other copy preserves the old epistemic structure for agents not in G . This ensures that they do not receive any information about the private announcement that took place.

Whereas public announcements preserve the structures of EAL models—the plausibility relations are well-founded, locally connected preorders—this is not true for private announcements. To ensure that the event takes place privately, the agents $b \notin G$ do not have access to e_ϕ in Definition 12, i.e. there is no reflexive relation for these agents at e_ϕ . This means that, in the product update, the worlds ‘created’ with e_ϕ do not have reflexive relations for $b \notin G$. As a consequence, the resulting models are no longer EAL models. To prevent this, typically event models are required to be such that they preserve the properties of the model structures [3]. Then it holds that applying an event to a EAL model yields a EAL model. In the situation here, this requirement boils down to requiring the relations R_a in event models to satisfy the same properties as the plausibility relations \leq_a in DEL models [3].

Of course, this restriction constraints the type of event models that can be applied. Another way to deal with this would be to weaken the logic. Later in this paper it is explained how weak reflexivity can be another way out.

Event models as described here have been generalized in [8] to accommodate factual change by introducing *postconditions* that define the new valuation function in the product update.

4 Agent awareness and unawareness

Dynamic Epistemic Logic is a rich framework for analyzing epistemic and doxastic changes under dynamic actions. However, there is a category of multi-agent systems for which DEL may be considered to fall short: *dynamic* and *open* multi-agent systems. In such systems, agents are typically required to continuously adapt to their environment and to interact while preserving heterogeneity in their knowledge representations [2, 20, 21].

There are two problems when it comes to using DEL to model dynamic and open multi-agent systems: (1) agents cannot use their *own* vocabularies in DEL,

and (2) agents cannot *extend* (nor shrink) their vocabularies. The reason for this is what we call public signature awareness: the agents use the same, fixed vocabulary. This means that agents are *aware* of all the propositions used by other agents, now or in the future, and that learning new information is limited to eliminating or re-organizing the possibilities with respect to these propositions.

The problem of public signature awareness is hidden in the structure of DEL: valuation functions are total functions and accessibility relations are reflexive. Therefore, in order to eliminate public signature awareness and to faithfully capture dynamic and open multi-agent systems, the structure of the models needs to be adjusted. For that purpose, total valuations are replaced by *partial valuations* and reflexivity by *weak reflexivity*. First, we motivate both before defining them precisely.

Partial valuation functions

By using total valuation functions, the vocabularies of agents become equivalent to the set of all propositions, P . This is because they enforce each proposition to be evaluated to 0 or 1 at each world, therefore enforcing agents to be aware of all of them. With partial valuation functions, on the contrary, propositions may either be true, false, or a third option: *undefined*. The latter occurs when the propositions are not interpreted by the valuation function at a certain world, i.e. they do not belong to the domain $Dom(V_w)$ of the valuation function V at some world w . By assigning a different partial valuation function to each world, we can specify both which propositions are evaluated (those belonging to the domain) and how they are evaluated (true or false), enabling agents to use different signatures, denoted by $sig(a) \subseteq P$ for agent a , to represent their knowledge and beliefs. This is achieved by defining the signature of agent a as the domain of the worlds w she can reach, i.e. $sig(a) = Dom(V_w)$. Later we will define natural restrictions on the domain so that the signature of agents is constant over the worlds they can reach and $sig(a) = Dom(V_w)$ is a proper definition.

Weakly reflexive relations

However, even with partial valuation functions, as long as reflexivity is satisfied, agents still share the same signature. This is because reflexive relations cause agents to have access to at least one common world w (such that $\langle \mathcal{M}, w \rangle$ is the pointed model to which formulas are evaluated) so that $sig(a) = Dom(V_w) = sig(b)$. Hence, we also need to drop reflexivity as an assumption on the model structures. Without reflexivity, two agents a and b can have access to different worlds v and u from w , respectively, such that $Dom(V_v) = sig(a)$ and $Dom(V_u) = sig(b)$ and neither $wR_a w$ nor $wR_b w$. In order to preserve consistency of agents' knowledge and beliefs, reflexivity is asserted to hold at v and u . That means that reflexivity is not satisfied globally, but locally at the worlds that determine the signatures (propositions defined) and the awareness (which worlds satisfy reflexivity) of the agents.

We call our alternative definition of reflexivity *weak reflexivity*. It requires that, whenever reflexivity holds at a world w for an agent a , it will continue to hold at any other world v agent a can access from w . As we will see later, as a consequence

the usual properties of knowledge and beliefs, and in particular factivity of knowledge, still hold within the *aware cells* of agents: those worlds in which reflexivity holds.

4.1 A definition of awareness

The accessibility relations indicate which worlds the agents are aware of—those where reflexivity is satisfied—and the partial valuation functions determine which propositions the agents are aware of—those propositions belonging to the domain of these worlds. This enables agents to use their own signatures to represent their knowledge and beliefs, and to adapt or extend it.

As a consequence of using partial valuation functions, lack of truth and falsity do no longer coincide. Instead, there is a third option: propositions may be *undefined*. Whenever this happens for a proposition p at a world accessible by an agent, this agent is said to be *unaware* of p .

Unawareness is different from uncertainty, which is also called *ignorance*. The latter occurs when agents have no information about the truth value of a proposition, i.e. they do not know nor believe the proposition, nor its truth value, whereas unawareness occurs when agents do not consider a proposition at all, i.e. the proposition is undefined in the worlds accessible by that agent. Uncertain agents have access to at least one world in which p is true and at least one world in which p is false, which are considered equally plausible, and unaware agents do not have access to any p -world nor $\neg p$ -world (Figure 1).




Fig. 1 Agent a is certain about p ($\models K_a p$, left), uncertain about p ($\models K_a(p \vee \neg p)$, middle) and unaware of p (right). In all cases, agent a knows q .

Unawareness of a proposition means that agents do not consider the proposition. Awareness of a proposition means that agents do consider the proposition, and hence also its truth value of it. Therefore, awareness of a proposition implies that agents are at least uncertain about the proposition.

The signature of an agent, also more generally called the awareness of an agent, is defined as the propositions evaluated to 0 or 1 in the worlds that she can access such that reflexivity holds.

In Figure 1, agent a is aware at w, v, w', v' and w'' , the awareness of agent a is $\{p, q\}$ in the models on the left and in the middle, and $\{q\}$ in the model on the right. Note that, like before, we write, inside a world w , p to denote that $V_w(p) = 1$ and \bar{p} to denote that $V_w(p) = 0$, but now, when we do not write p nor \bar{p} , it means that $p \notin \text{Dom}(V_w)$.

Definition 13 (Agent Awareness) If W is a non-empty set of worlds, P a countable, non-empty set of propositions, \mathcal{A} a finite, non-empty set of agents,

$\{V_w\}_{w \in W}$ a set of partial valuation functions $V_w : P \rightarrow \{0, 1\}$, and a an agent with accessibility relation $R_a \subseteq W \times W$, then a is said to be *aware* at w if and only if $wR_a w$ and the *awareness* (or signature, or vocabulary) of agent a at w is defined as:

$$AW_a(w) = \bigcup_{\{v \in W \mid w(R_a \cup R_a^{-1})v \wedge vR_a w\}} Dom(V_v) \quad (4)$$

Whenever w is clear from the context in a pointed model $\langle \mathcal{M}, w \rangle$, we will also use $AW_a(\langle \mathcal{M}, w \rangle)$, $AW_a(\mathcal{M})$ or simply AW_a for $AW_a(w)$ to denote the awareness of agent a in the pointed model.




Fig. 2 Two agents, a and b , with different, disjoint signatures, $\{p\}$ and $\{q\}$, respectively. Agents are not aware of the same propositions.

In Figure 2, the awareness of agent a is $\{p\}$ and the awareness of agent b is $\{q\}$.

4.2 Properties of awareness

Partial valuation functions and a weaker form of reflexivity are necessary for agents to use different signatures to represent their knowledge and beliefs. Here we discuss the properties of awareness that we require.

It is a natural assumption to let awareness of agents be constant over their accessibility relations. This is to ensure that agents do not consider it possible that their awareness is different from what they are actually aware of, i.e. what one can consider is limited to one's awareness. This means that the awareness of agent is constant over their considerations. Similar properties for awareness were already motivated: in [22], awareness is assumed to not be able to decrease under evolution (i.e. there is no “forgetting”) and in [25] awareness is considered constant for all the worlds the agent has access to. Compared to these works, the properties of awareness we introduce here are not different, but the framework for awareness itself is (using partial valuation functions and weakly reflexive relations).

Take note that requiring that awareness is constant over accessibility does not imply that it is fixed: agents can extend their awareness through model-changing

dynamic upgrades that is not restricted to the set of propositions P , which we will define in Section 6 on raising awareness.

Letting agent awareness be constant over their accessibility comes twofold:

- whenever there is a reflexive relation for an agent a from a world w to w , then for any v that is also accessible from w for a , there is a reflexive relation from v to v (*weak reflexivity*), and
- the domains of two worlds v and u that are both accessible by the same agent from a single world w are equal (*consideration consistency*).

Weak reflexivity ensures that if an agent is aware at a world w , she will remain aware at all the worlds she can reach from w , see Figure 3.

Definition 14 (Weak Reflexivity) Let W be a non-empty set of worlds and let $a \in \mathcal{A}$ be an agent with a relation $R_a \subseteq W \times W$. Then R_a is called *weakly reflexive* if $\forall w, v \in W$:

$$wR_a w \wedge wR_a v \Rightarrow vR_a v \quad (5)$$




Fig. 3 An illustration of weak reflexivity: if the black arrows hold for a certain agent a , then there must also be the red arrow for a .

Consideration consistency ensures that the awareness of agents, i.e. the propositions that are defined in a world with reflexivity, is constant over accessibility. Figure 4 illustrates this property.

Definition 15 (Consideration Consistency) Let W be a non-empty set of worlds, P be a countable, non-empty set of propositions, \mathcal{A} be a finite, non-empty set of agents, $\{V_w\}_{w \in W}$ be a set of partial valuation functions, and let $a \in \mathcal{A}$ be an agent with a relation $R_a \subseteq W \times W$. Then $\{V_w\}_{w \in W}$ satisfies *consideration consistency* if $\forall w, v, u \in W$:

$$wR_a v \wedge wR_a u \Rightarrow \text{Dom}(V_v) = \text{Dom}(V_u) \quad (6)$$

Consideration consistency does not only apply to worlds w, v, u such that $wR_a v$ and $wR_a u$ (hence $\text{Dom}(V_v) = \text{Dom}(V_u)$), but also stipulates that if $wR_a w$ and $wR_a v$ then $\text{Dom}(V_w) = \text{Dom}(V_v)$. Therefore when combining consideration consistency with weak reflexivity, it enforces that agents are consistent in their considerations: if an agent a considers a proposition p (or its negation) to be true at a world w (i.e. $p \in \text{Dom}(V_w)$ and $wR_a w$), she considers p to have a truth value at every world v she can reach via R_a from w (i.e. also $p \in \text{Dom}(V_v)$ and $vR_a v$), and vice-versa. This is independent from the actual truth value of p —it only requires that p is *assigned* a truth value 0 or 1.




Fig. 4 An illustration of consideration consistency: if the arrows hold for a certain agent a , and the domains are $\text{Dom}(V_v)$ and $\text{Dom}(V_u)$, then these domains are equal, here drawn to be P' .

Proposition 1 (Awareness is constant over accessibility) *Let W be a set of worlds, P be a set of propositions, $\{V_w\}_{w \in W}$ a set of partial valuation functions satisfying consideration consistency, and let a be an agent with a relation $R_a \subseteq W \times W$ that is weakly reflexive. It holds that $\forall w \in W$: if $wR_a w$ then $\forall v \in W$ such that $wR_a^* v$, we have $vR_a v$ and $AW_a(v) = AW_a(w)$, where R_a^* is the transitive closure of R_a .*

Proof Follows directly by Definitions 14 and 15 for weak reflexivity and consideration consistency.

Lastly, we require that agents cannot reason about what other agents know or believe when it involves a proposition they are not aware of themselves. For example, in Figure 2, if $\text{Dom}(V_t)$ included q , it would allow agent a to know that b knows q , whereas a is unaware of q . To avoid this, we require a last condition called *specification*:

- the domain of the valuation function cannot increase over accessibility (*specification*).

Specification avoids the situation in which, for example, $K_a K_b p$ but $p \notin AW_a$. This will be further discussed in Section 4.3. Note that when $wR_a w$ and $wR_a v$, by consideration consistency, $\text{Dom}(V_v) = \text{Dom}(V_w)$ and therefore $\text{Dom}(V_v) \subseteq \text{Dom}(V_w)$. Figure 5 illustrates specification.

Definition 16 (Specification) Let W be a non-empty set of worlds, P be a countable, non-empty set of propositions, \mathcal{A} be a finite, non-empty set of agents, $\{V_w\}_{w \in W}$ be a set of partial valuation functions, and let a be an agent with a relation $R_a \subseteq W \times W$. Then $\{V_w\}_{w \in W}$ satisfies *specification* if $\forall w, v \in W$:

$$wR_a v \Rightarrow \text{Dom}(V_v) \subseteq \text{Dom}(V_w) \quad (7)$$

The three properties—weak reflexivity, consideration consistency and specification—ensure that agents are consistent in their considerations, while allowing models in which agents are aware of different signatures, or vocabularies. For example, the model shown in Figure 2 satisfies weak reflexivity, consideration consistency and specification, yet the awareness of agents is different: $\{p\}$ for agent a and $\{q\}$ for agent b .

Lastly, it may be clear that if for two agents a and b reflexivity is satisfied at w , their awareness is the same, i.e. $AW_a(w) = AW_b(w)$, by consideration consistency.




Fig. 5 An illustration of specification: if the arrow from w to v holds for a certain agent a , and the domains are $Dom(V_w) = P_1$ and $Dom(V_v) = P_2$, then $P_2 \subseteq P_1$.

4.3 Awareness for knowledge and beliefs

We define the knowledge and beliefs of agents as usual with respect to the accessibility relations. However, since reflexivity is not satisfied globally, there is an additional requirement: reflexivity needs to hold. Thus knowledge is defined as everything that is true in all accessible worlds for an agent in which reflexivity is satisfied for that agent, and belief is defined as everything that is true in the maximal worlds with respect to the accessibility relation for an agent in which reflexivity is satisfied. These notions can also be captured by *epistemic* and *doxastic* relations, which are deduced from the accessibility relations. In particular, they are defined with respect to the *aware cell* of an agent. This aware cell consists of all the worlds satisfying reflexivity accessible by an agent, see Figure 6 for an example.

Definition 17 (Aware Cell) Let W be a non-empty set of worlds and a finite, non-empty set \mathcal{A} of agents and let $R_a \subseteq W \times W$ be an accessibility relation for agent a that is well-founded, locally connected, weakly reflexive and transitive. Then the *aware cell* of agent a at world $w \in W$, denoted by $\|w\|_a$, is the set of worlds that are accessible via the transitive closure of R_a and R_a^{-1} and in which reflexivity is satisfied:

$$\|w\|_a = \{v \in W \mid w(R_a \cup R_a^{-1})^*v \text{ and } vR_av\} \quad (8)$$




Fig. 6 The aware cells are as follows: $\|w\|_a = \|v\|_a = \{v\}$ and $\|w'\|_a = \|v'\|_a = \{w', v'\}$ (marked by dotted boxes).

Within the aware cells, the awareness of agents is constant: $\forall w, u, v \in W$, $\forall a \in \mathcal{A}$: if $u, v \in \|w\|_a$ then $AW_a(v) = Dom(V_v) = Dom(V_u) = AW_a(u)$. Therefore we also simply use $Dom(V_w)$ to denote the awareness of agent a at world w .

The epistemic and doxastic relations are defined with respect to the aware cell of an agent.

Definition 18 (Epistemic and Doxastic Relations) Let \mathcal{A} be a finite, non-empty set of agents, W be a non-empty set of worlds and let $a \in \mathcal{A}$ be an agent with

accessibility relation $R_a \subseteq W \times W$ that is well-founded, locally connected, weakly reflexive and transitive. Then the epistemic (\sim_a) and doxastic (\rightarrow_a) relations are defined as follows:

$$w \sim_a v \text{ iff } v \in ||w||_a \quad (9)$$

$$w \rightarrow_a v \text{ iff } v \in \text{Max}_{R_a} ||w||_a \quad (10)$$

Within aware cells of an agent a , \sim_a and \rightarrow_a satisfy the usual properties for epistemic and doxastic relations. Therefore, within the aware cells of agents, knowledge and beliefs satisfy the usual axioms as for DEL: $S5$ and $KD45$, respectively.

Proposition 2 *Let W be a set of worlds and a an agent with accessibility relation $R_a \subseteq W \times W$ that is well-founded, locally connected, weakly reflexive and transitive. Then within $||w||_a$, \sim_a is reflexive, transitive and symmetric and \rightarrow_a is transitive, serial and Euclidean.*

Proof The only difference with the plausibility relation \leq_a for DEL (Definition 2) is that R_a is not reflexive, but weakly reflexive. By Definition 17 of aware cells, we have that $\forall w \in W$ and $\forall v \in ||w||_a: vR_av$. Hence, within $||w||_a$, the relation R_a is a well-founded locally connected pre-order and it therefore follows that \sim_a and \rightarrow_a satisfy the same properties as the epistemic and doxastic relations for DEL [19]: \sim_a is reflexive, transitive and symmetric and \rightarrow_a is transitive, serial and Euclidean.

However, outside the aware cells, knowledge is not necessarily factive, i.e. the T axiom is no longer valid. Instead, it is only valid within the aware cells of agents. This behavior is due to weak reflexivity: reflexivity holds only in the aware cells of agents. As a consequence, the knowledge operators K_a are intermediate operators between $S5$ and $KD45$: K_a satisfies $S5$ within the aware cell of agent a , but satisfies $KD45$ outside the aware cell. In practice, such a knowledge operator may be thought of as “subjective knowledge”: it takes the perspective of the agent in question and what is known by her.

5 Partial Dynamic Epistemic Logic

We are now ready to define our logic, Partial Dynamic Epistemic Logic (ParDEL). The syntax of ParDEL is equivalent to that of DEL.

Definition 19 (Syntax of ParDEL) Given a non-empty set P of propositions and a finite non-empty set \mathcal{A} of agents, the *syntax*, $\mathcal{L}_{\text{ParDEL}}$, of (multi-agent) Partial Dynamic Epistemic Logic (ParDEL) is defined in the following way.

$$\phi ::= p \mid \phi \wedge \psi \mid \neg\phi \mid K_a\phi \mid B_a\phi \mid [\dagger\phi]\psi$$

where $p \in P$ is a proposition, K_a and B_a are the knowledge and belief operators for each agent $a \in \mathcal{A}$, and $\dagger\phi$ with $\dagger \in \{!, \uparrow, \updownarrow\}$ the dynamic upgrades: announcements, radical and conversative upgrades.

ParDEL frames are DEL frames but satisfy weak reflexivity instead of reflexivity in order to allow agents to use different signatures to represent their knowledge and beliefs.

Definition 20 (ParDEL Frames) Given a finite non-empty set \mathcal{A} of agents, a *frame* of (multi-agent) ParDEL is a pair $\mathfrak{F} = \langle W, (R_a)_{a \in \mathcal{A}} \rangle$ where

- W is a non-empty set of worlds, and
- $(R_a)_{a \in \mathcal{A}} : \mathcal{A} \rightarrow \mathcal{P}(W \times W)$ are the accessibility relations on W , one for each agent, that are well-founded, locally connected, weakly reflexive and transitive.

A ParDEL model is a ParDEL frame equipped with a partial valuation function satisfying consideration consistency and specification.

Definition 21 (ParDEL Models) Given a non-empty set P of propositions and a finite non-empty set \mathcal{A} of agents, a *model* of (multi-agent) ParDEL is a pair $\mathcal{M} = \langle \mathfrak{F}, V \rangle$ where

- $\mathfrak{F} = \langle W, (R_a)_{a \in \mathcal{A}} \rangle$ is a ParDEL frame, and
- $V : W \rightarrow (P \rightarrow \{0, 1\})$ is a *partial valuation function* that assigns to each world $w \in W$ a partial function $V_w : P \rightarrow \{0, 1\}$ satisfying consideration consistency and specification.

A *pointed ParDEL model* is a pair $\langle \mathcal{M}, w \rangle$, where $w \in W$.

Satisfiability for ParDEL is considered with respect to a pointed model $\langle \mathcal{M}, w \rangle$ which associates a ParDEL model \mathcal{M} with a world $w \in W$. Since lack of truth and falsity do not coincide when considering partial valuation functions, two relations are specified: one for *verification* (\models) and one for *falsification* ($\models\!\!\!\neq$).

Satisfiability for ParDEL is based on satisfiability for partial semantics in [32] and extends it in three ways: (1) by introducing a novel connection between partial valuations and agent awareness (Section 4), (2) by considering knowledge and beliefs of agents as truth in all accessible worlds and most plausible worlds in which *reflexivity is satisfied* (Section 4.3) and (3) by considering a different clause for the falsification of a conjunction: a conjunction is false whenever at least one of the conjuncts is false, as usual, but in addition both conjuncts belong to the domain of the valuation function. In particular, we require (3) in order to prevent agents from gaining the knowledge that a conjunction is false (or similarly that a disjunction is true) whenever they know that one of the conjuncts is false (or one of the disjuncts is true), but are unaware of the other conjunct (or disjunct). This enforces that the knowledge and beliefs of agents are constrained to their awareness.

Definition 22 (Satisfiability for ParDEL) Satisfiability for ParDEL extends that of DEL (Definition 5) and is defined as:

$$\begin{array}{ll}
\mathcal{M}, w \models p & \text{iff } V_w(p) = 1 \\
\mathcal{M}, w \models \phi \wedge \psi & \text{iff } \mathcal{M}, w \models \phi \text{ and } \mathcal{M}, w \models \psi \\
\mathcal{M}, w \models \neg\phi & \text{iff } \mathcal{M}, w \models\!\!\!\neq \phi \\
\mathcal{M}, w \models K_a\phi & \text{iff } \forall v \text{ s.t. } w \sim_a v : \mathcal{M}, v \models \phi \\
\mathcal{M}, w \models B_a\phi & \text{iff } \forall v \text{ s.t. } w \rightarrow_a v : \mathcal{M}, v \models \phi \\
\mathcal{M}, w \models [\dagger\phi]\psi & \text{iff } \mathcal{M}^{\dagger\phi}, w \models \psi
\end{array}$$

for verification (\models) and for falsification (\models):

$$\begin{array}{ll}
\mathcal{M}, w \models p & \text{iff } V_w(p) = 0 \\
\mathcal{M}, w \models \phi \wedge \psi & \text{iff } \mathcal{M}, w \models \phi \text{ and } \mathcal{M}, w \models \psi, \\
& \text{or } \mathcal{M}, w \models \phi \text{ and } \mathcal{M}, w \models \psi, \\
& \text{or } \mathcal{M}, w \models \phi \text{ and } \mathcal{M}, w \models \psi \\
\mathcal{M}, w \models \neg\phi & \text{iff } \mathcal{M}, w \models \phi \\
\mathcal{M}, w \models K_a\phi & \text{iff } \exists v \text{ s.t. } w \sim_a v : \mathcal{M}, v \models \phi \\
\mathcal{M}, w \models B_a\phi & \text{iff } \exists v \text{ s.t. } w \rightarrow_a v : \mathcal{M}, v \models \phi \\
\mathcal{M}, w \models [\dagger\phi]\psi & \text{iff } \mathcal{M}^{\dagger\phi}, w \models \psi
\end{array}$$

with $\dagger \in \{\downarrow, \uparrow, \uparrow\}$.

As usual, a set of formulas is said to be inconsistent if there does not exist a pointed model verifying it. In the following, we say that a formula ϕ is a consequence of a set of formulas Γ (written $\Gamma \models \phi$) if every pointed model $\langle \mathcal{M}, w \rangle$ verifying all formulas of Γ , also verifies ϕ .

Whenever a proposition p does not belong to the domain of the valuation function at a world w , i.e. $p \notin \text{Dom}(V_w)$, we have indeed that $\mathcal{M}, w \not\models p$ and $\mathcal{M}, w \not\models p$. Hence, p is neither true nor false.

As disjunctions and implications can be defined using conjunctions and negations, the falsification clause for conjunctions also affects their satisfiability. In particular, disjunctions are only true if both disjuncts are defined and at least one of them is true. This ensures that agents only know or believe a disjunction if they are aware of both disjuncts.

Example 3 Since $\phi \vee \psi$ is the abbreviation for $\neg(\neg\phi \wedge \neg\psi)$:

$$\begin{array}{ll}
\mathcal{M}, w \models \phi \vee \psi & \text{iff } \mathcal{M}, w \models \phi \text{ and } \mathcal{M}, w \models \psi, \\
& \text{or } \mathcal{M}, w \models \phi \text{ and } \mathcal{M}, w \models \psi, \\
& \text{or } \mathcal{M}, w \models \phi \text{ and } \mathcal{M}, w \models \psi \\
\mathcal{M}, w \models \phi \vee \psi & \text{iff } \mathcal{M}, w \models \phi \text{ and } \mathcal{M}, w \models \psi
\end{array}$$

There is a link between intuitionism [13] and awareness, in particular the way satisfiability is defined for ParDEL. Example 3 shows that the truth of a disjunction is defined constructively: $\phi \vee \psi$ is true as long as *both* disjuncts are defined, of which at least one is true. That means that satisfiability for ParDEL follows that of Weak Kleene Logic [33]. In this sense, we may say that awareness is constructive. Other than that, however, awareness and intuitionism are rather orthogonal: intuitionism drops the law of the excluded middle, dissociating ϕ from $\neg\phi$, whereas awareness dissociates being aware of ϕ from knowing the truth value of ϕ and $\neg\phi$.

6 The dynamics of raising awareness

Partial valuation functions allow agents to use different signatures to represent their knowledge. We have seen how these can be used to describe awareness and unawareness of agents directly in the structure of the models when we replace

reflexivity by weak reflexivity. We have also seen some properties of awareness that are natural to require: consideration consistency and specification. Here, we discuss some dynamics for awareness. In particular, the dynamics of raising awareness. Because even though awareness of agents is constant over their accessibility, partial valuations enable to define a model-changing dynamic upgrade allowing agents to extend their awareness.

We extend the syntax of ParDEL with a dynamic modality $+p$ for raising the awareness of a proposition p . We call the logic obtained *Partial Dynamic Epistemic Logic with raising awareness* (ParDEL+).

Definition 23 (Syntax of ParDEL+) Given a non-empty set P of propositions and a finite non-empty set \mathcal{A} of agents, the *syntax*, $\mathcal{L}_{ParDEL+}$, of (multi-agent) Partial Dynamic Epistemic Logic with raising awareness (ParDEL+) is defined in the following way.

$$\phi ::= p \mid \phi \wedge \psi \mid \neg\phi \mid K_a\phi \mid B_a\phi \mid [\dagger\phi]\psi \mid [+p]\phi$$

where $p \in P$ is a proposition, K_a and B_a are the knowledge and belief operators for each agent $a \in \mathcal{A}$, $\dagger\phi$ with $\dagger \in \{!, \uparrow, \uparrow\}$ the dynamic upgrades announcements, radical and conversative upgrades, and $+p$ is an operation for raising awareness.

Frames and models of ParDEL+ are equal to frames and models of ParDEL (Definitions 20 and 21), using accessibility relations that are weakly reflexive and partial valuation functions satisfying consideration consistency and specification.

When awareness was introduced, it was discussed how to interpret awareness of a proposition: as considering at least one truth value of the proposition. This means that awareness at least implies ignorance. Becoming ignorant is therefore the minimal way to raise awareness without disclosing truth.

To raise the awareness of a proposition p , we define a model transformation that extends the valuation function at each world in which p does not belong to the domain to include p . In order to solely raise awareness, i.e. to perform it without disclosing truth values, it does so by first duplicating each of these worlds, accessibility to and from duplicated worlds being preserved, such that p is made true in one world and false in the other. Other than the different valuation of p (and related sentences), the two duplicated worlds are indifferent, satisfying the same valuations and relations. After raising the awareness of p , agents come to consider equally plausible either that p is true or that p is false, see Figure 7. Therefore, raising the awareness of p transforms unaware agents into uncertain agents: agents who are aware of p but do not know whether p is true or false.

When raising the awareness of p , we categorize the worlds of the model into three: worlds in which p is true ($W|_p = \{w \in W \mid V_w(p) = 1\}$), worlds in which p is false ($W|_{\neg p} = \{w \in W \mid V_w(p) = 0\}$) and worlds in which p is undefined ($W \setminus (W|_p \cup W|_{\neg p})$). In this way, the worlds in which p is undefined can be identified to be duplicated by raising awareness.

Definition 24 (Raising Awareness (+p)) Let $\mathcal{M} = \langle W, (R_a)_{a \in \mathcal{A}}, V \rangle$ be a ParDEL+ model and let $p \in P$ be a proposition. Then $+p$ is a model transformer $+p : \mathcal{M} \mapsto \mathcal{M}^{+p}$ where \mathcal{M}^{+p} is the triple $\langle W^{+p}, (R_a^{+p})_{a \in \mathcal{A}}, V^{+p} \rangle$ defined by:

- $W^{+p} = W|_p \times \{1\} \cup W|_{\neg p} \times \{0\} \cup W \setminus (W|_p \cup W|_{\neg p}) \times \{0, 1\}$
- $\langle w, i \rangle R_a^{+p} \langle v, j \rangle$ iff $w R_a v$




Fig. 7 Raising the awareness of p , $+p$. The red dashed lines indicate how the world w from the model on the left is mapped to both $\langle w, 1 \rangle$ and $\langle w, 0 \rangle$ in the model on the right. p (\bar{p}) written inside a world w means that $V_w(p) = 1$ ($V_w(p) = 0$), and whenever neither p nor \bar{p} is written, it means that p does not belong to $\text{Dom}(V_w)$.

$$- V_{\langle w, i \rangle}^{+p}(q) = \begin{cases} V_w(q) & \text{if } q \neq p \\ i & \text{otherwise} \end{cases}$$

The new valuation function corresponds to the old one in the case that p is defined: $V_{\langle w, i \rangle}^{+p}(p) = i = V_w(p)$ because only in this case $\langle w, 1 \rangle \in W^{+p}$, where $i = 1$ if $V_w(p) = 1$ and $i = 0$ if $V_w(p) = 0$ (Figure 8). For worlds where p is undefined, i.e. $p \notin \text{Dom}(V_w)$, raising awareness of p maps w to both $\langle w, 1 \rangle$ and $\langle w, 0 \rangle$, in which p is made true and false, respectively.

Because worlds may be duplicated, satisfiability for $[+p]\phi$ is defined for all $\langle w, i \rangle \in W^{+p}$ with $i \in \{0, 1\}$ for verification, and for some $\langle w, i \rangle \in W^{+p}$ with $i \in \{0, 1\}$ for falsification.

Definition 25 (Satisfiability for ParDEL+) Satisfiability for ParDEL+ extends that of ParDEL (Definition 22) as follows:

$$\begin{aligned} \mathcal{M}, w \models [+p]\phi & \quad \text{iff } \forall \langle w, i \rangle \in W^{+p} \text{ s.t. } i \in \{0, 1\} : \mathcal{M}^{+p}, \langle w, i \rangle \models \phi \\ \mathcal{M}, w \models \neg [+p]\phi & \quad \text{iff } \exists \langle w, i \rangle \in W^{+p} \text{ s.t. } i \in \{0, 1\} : \mathcal{M}^{+p}, \langle w, i \rangle \models \neg \phi \end{aligned}$$

As we discussed before, in Figure 8, we can also see that knowledge is no longer factive, *except* within the aware cells of agents: it holds that $\langle w, 0 \rangle \models K_a p$ implies $\langle w, 0 \rangle \models p$ because $\langle w, 0 \rangle \in \|\langle w, 0 \rangle\|_a$, hence K_a is factive at $\langle w, 0 \rangle$. However, $\langle w, 0 \rangle \models K_b \neg K_a p$ holds, but $\langle w, 0 \rangle \not\models \neg K_a p$ because $\langle w, 0 \rangle \notin \|\langle w, 0 \rangle\|_b$. Therefore K_b is not factive at $\langle w, 0 \rangle$.

Raising awareness needs not to be constrained to the set of propositions P . We can use the same definition to raise the awareness of $q \notin P$ and extend P with q . In this case, q will be undefined at every world of the model, leaving each world to be duplicated by $+q$. This means that P does not need to be fixed in the initial setting.

6.1 Raising awareness with event models

Similar to announcements, we can capture the dynamics of raising awareness in an alternative way through event structures that specify how agents observe the




Fig. 8 Raising the awareness of a proposition p , $+p$. The red dashed lines indicate how the world w from the model on the left is mapped to $(w, 1)$ in the model on the right and v to both $(v, 0)$ and $(v, 1)$.

events. Just as event structures enable us to capture private announcements on Epistemic Action Logic (EAL), event structures for Partial Epistemic Action Logic (ParEAL), the logic combining the syntax of EAL with the semantics of ParDEL, enable us to define private raising awareness.

ParEAL follows the syntax of EAL (Definition 7) with modalities for *multi-pointed event models* $\langle \mathcal{E}, E' \rangle$.

Definition 26 (Syntax of ParEAL) Given a countable, non-empty set P of propositions and a finite, non-empty set \mathcal{A} of agents, the *syntax*, $\mathcal{L}_{\text{ParEAL}}$, of (multi-agent) Partial Epistemic Action Logic (ParEAL) is defined in the following way.

$$\phi ::= p \mid \phi \wedge \psi \mid \neg \phi \mid K_a \phi \mid B_a \phi \mid [\langle \mathcal{E}, E' \rangle] \phi$$

where $p \in P$ is a proposition, K_a and B_a are the knowledge and belief operators for each agent a , and $\langle \mathcal{E}, E' \rangle$ are multi-pointed event models.

The semantics of ParEAL is that of ParDEL and ParDEL+: frames and models of ParEAL are equivalent to ParDEL and ParDEL+ frames and models (Definitions 20 and 21), with accessibility relations satisfying weak reflexivity and partial valuations satisfying consideration consistency and specification.

Multi-pointed event models for ParEAL are based on event models for EAL (Definition 8), but with an added feature to specify the valuation of the proposition awareness is raised of: *postconditions* $post_e$ for event e .

Event models have been extended before with postconditions in [8] to accommodate for factual change, through a substitution function. Here, we consider a different definition of a postcondition. Postconditions for ParEAL event models are functions assigning to each event e a partial function $post_e : P \rightarrow \{0, 1\}$ such that this becomes the new valuation of p , if defined, and preserves the old valuation if $post_e(p)$ is undefined. This can likewise be used to define factual change, by letting $post_e(p)$ be equal to 0 or 1.

Definition 27 (Event Model with Postconditions) Let \mathcal{A} be a finite, non-empty set of agents. An *event model* for ParEAL is a quadruple $\mathcal{E} = \langle E, (R_a)_{a \in \mathcal{A}}, pre, post \rangle$ where

- E is a non-empty, finite set of *events*,
- $(R_a)_{a \in \mathcal{A}} \subseteq E \times E$ are the *accessibility relations* on E , one for each agent $a \in \mathcal{A}$,
- $pre : E \rightarrow \mathcal{L}_{ParDEL+}$ is a *precondition function* assigning to each event a formula ϕ , and
- $post : E \rightarrow (P \rightarrow \{0, 1\})$ is a *postcondition function* assigning to each event a partial function $post_e : P \rightarrow \{0, 1\}$.

A *pointed event model (with postconditions)* is a pair $\langle \mathcal{E}, e \rangle$ where \mathcal{E} is an event model with postconditions and $e \in E$.

A *multi-pointed event model (with postconditions)* is a pair $\langle \mathcal{E}, E' \rangle$ where \mathcal{E} is an event model with postconditions and $E' \subseteq E$.

We will also write pre_e for $pre(e)$ and $post_e(p)$ for $(post(e))(p)$.

Given an event model \mathcal{E} and a set of events $\{e_1, \dots, e_n\} \subseteq E$, the multi-pointed event model $\langle \mathcal{E}, \{e_1, \dots, e_n\} \rangle$ describes the set of pointed event models $\langle \mathcal{E}, e \rangle$ with $e \in \{e_1, \dots, e_n\}$. When drawing a multi-pointed event model $\langle \mathcal{E}, \{e_1, \dots, e_n\} \rangle$ for ParEAL, events are again drawn as squares to distinguish them from ParDEL models and all $e \in \{e_1, \dots, e_n\}$ are double-squared to emphasize the points of reference.

Like for DEL, the product update of a ParEAL model \mathcal{M} and an event model \mathcal{E} determines what happens if the event takes place. In this product update, the preconditions again specify how to select the worlds in the product update. However, different to DEL, in ParEAL this selection takes place on the basis of ‘not falsifying the precondition’, instead of requiring that the precondition is verified. This is because lack of truth and falsity do not coincide on ParEAL. As a result, a precondition p selects the worlds in which either p is true, or p is undefined. That is, the worlds in which the precondition is undefined will remain too. This ensures that we can capture raising awareness in a natural way, as we will see.

Additionally, postconditions are used to specify the valuations that change in the product update. When the postcondition of a proposition is 1, this proposition becomes true, when it is 0, it becomes false, and when it is \perp , it becomes undefined. Lastly, when the postcondition for a proposition is undefined, the old valuation is preserved.

Definition 28 (Product Update for ParEAL) Let $\mathcal{M} = \langle W, (R_a^{\mathcal{M}})_{a \in \mathcal{A}}, V \rangle$ be a ParEAL model and $\mathcal{E} = \langle E, (R_a^{\mathcal{E}})_{a \in \mathcal{A}}, pre, post \rangle$ be an event model. Their *product update*, denoted by $\mathcal{M} \otimes \mathcal{E}$, is the triple $\langle W^{\mathcal{M} \otimes \mathcal{E}}, (R_a^{\mathcal{M} \otimes \mathcal{E}})_{a \in \mathcal{A}}, V^{\mathcal{M} \otimes \mathcal{E}} \rangle$ defined by:

- $W^{\mathcal{M} \otimes \mathcal{E}} = \{ \langle w, e \rangle \in W \times E \mid \mathcal{M}, w \not\models pre_e \}$
- $\langle w, e \rangle R_a^{\mathcal{M} \otimes \mathcal{E}} \langle w', e' \rangle$ iff $\langle w, e \rangle, \langle w', e' \rangle \in W^{\mathcal{M} \otimes \mathcal{E}}$, $w R_a^{\mathcal{M}} w'$ and $e R_a^{\mathcal{E}} e'$
- $V_{\langle w, e \rangle}^{\mathcal{M} \otimes \mathcal{E}}(p) = \begin{cases} post_e(p) & \text{if } post_e(p) = 1 \text{ or } post_e(p) = 0 \\ V_w(p) & \text{otherwise} \end{cases}$

In the following we also refer to the events e such that $\mathcal{M}, w \not\models pre(e)$ as the events that *can be applied* to w . It follows from the definition that whenever no postcondition is defined, this means that the old valuation is preserved completely.

The event model for raising awareness is as follows.

Definition 29 (Event Model for Raising Awareness) The multi-pointed event model for *raising awareness* of a proposition $p \in P$ is the pair $\langle \mathcal{E}_{+p}, \{e_p, e_{\bar{p}}\} \rangle$ where $\mathcal{E}_{+p} = \langle E_{+p}, (R_a)_{a \in \mathcal{A}}, pre, post \rangle$, with $E_{+p} = \{e_p, e_{\bar{p}}\}$, $R_a = \{e_p, e_{\bar{p}}\} \times \{e_p, e_{\bar{p}}\}$ and the pre- and postconditions defined as follows (see Figure 9):

- $pre_{e_p} = p, post_{e_p}(p) = 1$
- $pre_{e_{\bar{p}}} = \neg p, post_{e_{\bar{p}}}(p) = 0$

We also use E_{+p}^* to denote the points of reference $\{e_p, e_{\bar{p}}\}$.

Because the pre- and postconditions ‘coincide’ (that is, the valuation defined by the postcondition verifies the precondition), we can draw the event model as follows, where written p inside an event e means that $pre_e = p$ and $post_e(p) = 1$ and written $\neg p$ means that $pre_e = \neg p$ and $post_e(p) = 0$:




Fig. 9 The event model \mathcal{E}_{+p} for raising awareness of $p, +p$.

The event model for raising awareness then duplicates all the worlds of a model in which p was undefined, making p in one world and false in the other, while it preserves the other worlds. This is because worlds w such that $V_w(p)$ is undefined do not falsify p nor $\neg p$, hence $\mathcal{W}, w \not\models pre_{e_p}$ and $\mathcal{W}, w \not\models pre_{e_{\bar{p}}}$ and both events e_p and $e_{\bar{p}}$ can be applied to w . For an example of the application of \mathcal{E}_{+p} to a model, see Figure 10.




Fig. 10 The event \mathcal{E}_{+p} applied to the epistemic model on the left.

Satisfiability is defined analogously as for EAL (Definition 10), but now with respect to multi-pointed event models. The multiple points of reference are used to enforce that, whenever p was initially undefined at w of $\langle \mathcal{M}, w \rangle$ and awareness of p is raised, so that it is duplicated in two worlds $\langle w, e_p \rangle$ and $\langle w, e_{\bar{p}} \rangle$, both these duplicated worlds must make ϕ true in order for $[+p]\phi$ to hold at w . In this way, satisfiability for the raising awareness event model is equivalent to the clause for $+p$ in Definition 25.

Definition 30 (Satisfiability for Multi-pointed Event Models) Given a multi-pointed event model $\langle \mathcal{E}, E' \rangle$, *satisfiability* for ParEAL extends satisfiability of ParDEL (Definition 22), replacing the clause for $[\dagger\phi]\psi$ by:

$$\begin{aligned} \mathcal{M}, w \models [\langle \mathcal{E}, E' \rangle] \phi \text{ iff } & \forall e \in E' : \mathcal{M}, w \not\models \text{pre}(e) : \mathcal{M} \otimes \mathcal{E}, \langle w, e \rangle \models \phi \\ \mathcal{M}, w \models \langle \mathcal{E}, E' \rangle \phi \text{ iff } & \exists e \in E' : \mathcal{M}, w \not\models \text{pre}(e) : \mathcal{M} \otimes \mathcal{E}, \langle w, e \rangle \models \phi \end{aligned}$$

Figures 8 and 10 illustrate that $+p$ as defined as a model transformer in Definition 24 and the multi-pointed event model \mathcal{E}_{+p} lead to the same model.

Proposition 3 *For any ParDEL+/ParEAL model \mathcal{M} , \mathcal{M}^{+p} is bisimilar to $\mathcal{M} \otimes \mathcal{E}_{+p}$.*

Proof \mathcal{M}^{+p} is \mathcal{M} with every world $w \in \mathcal{M}$ such that $p \notin \text{Dom}(V_w)$ duplicated into $\langle w, 0 \rangle$ and $\langle w, 1 \rangle$ making p false and true, respectively, while the relations and valuations for other propositions remain the same. This is exactly the same as $\mathcal{M} \otimes \mathcal{E}_{+p}$, by replacing 1 by e_p and 0 by $e_{\bar{p}}$. I.e. the same valuations and the same relations hold. Therefore Z defined as $\forall \langle w, i \rangle \in W^{+p} : \langle w, 0 \rangle Z \langle w, e_{\bar{p}} \rangle$ and $\langle w, 1 \rangle Z \langle w, e_p \rangle$ is a bisimulation.

Furthermore, the semantics is equivalent for $+p$ and $\langle \mathcal{E}_{+p}, \{e_p, e_{\bar{p}}\} \rangle$. Therefore, we can use these two interchangeably.

Proposition 4 *For any ParDEL+/ParEAL model \mathcal{M} , any world $w \in \mathcal{M}$, any proposition $p \in P$ and any formula ϕ , it holds that:*

$$\begin{aligned} \mathcal{M}, w \models [+p] \phi \text{ iff } \mathcal{M}, w \models [\langle \mathcal{E}_{+p}, \{e_p, e_{\bar{p}}\} \rangle] \phi \\ \mathcal{M}, w \models \langle [+p] \phi \text{ iff } \mathcal{M}, w \models [\langle \mathcal{E}_{+p}, \{e_p, e_{\bar{p}}\} \rangle] \phi \end{aligned}$$

Proof We prove the case of verification, the case of falsification is analogous. It holds that $\mathcal{M}, w \models [+p] \phi$ iff $\forall \langle w, i \rangle \in W^{+p}$ with $i \in \{0, 1\} : \mathcal{M}^{+p}, \langle w, i \rangle \models \phi$. But since \mathcal{M}^{+p} and $\mathcal{M} \otimes \mathcal{E}_{+p}$ are bisimilar (Proposition 3), they can be exchanged so that, by renaming $\langle w, i \rangle$ with events $e_p, e_{\bar{p}}$, $\mathcal{M}, w \models [+p] \phi$ iff $\forall \langle w, e \rangle \in W^{\mathcal{M} \otimes \mathcal{E}}$ with $e \in \{e_p, e_{\bar{p}}\} : \mathcal{M} \otimes \mathcal{E}, \langle w, e \rangle \models \phi$. Therefore, $\mathcal{M}, w \models [+p] \phi$ iff $\mathcal{M}, w \models [\langle \mathcal{E}_{+p}, \{e_p, e_{\bar{p}}\} \rangle] \phi$.

6.2 Raising private awareness

When we discussed event models for DEL (see Section 3.1), we explained how these models enable us to study more complex epistemic upgrades such as private announcements, by computing the product update [3]. Analogous to private announcements, in this section we consider the dynamics of raising *private* awareness: a group G of agents raises their awareness of a proposition p , but this occurs in full privacy to the other agents ($\mathcal{A} \setminus G$). This means that the other agents do not observe the upgrade, they do not even consider that it occurred.

We can capture raising private awareness with an event model that looks similar to the event model for public awareness (Definition 29), but instead of two events, one with precondition p and another with precondition $\neg p$, a third event needs to be considered: an event with precondition \top , analogous to e_{\top} in the event model for private announcements (see Definition 12). Like for public raising awareness, the first two events duplicate the worlds in which p is undefined, making it true in one and false in the other, but now the third event ensures that the other agents do not observe this to occur.

Definition 31 (Event Model for Private Raising Awareness) The multi-pointed event model for *private raising awareness* of a proposition p amongst a group $G \subseteq \mathcal{A}$ of agents is $\langle \mathcal{E}_{+Gp}, \{e_p, e_{\bar{p}}\} \rangle$ where $\mathcal{E}_{+Gp} = \langle E_{+Gp}(R_a)_{a \in \mathcal{A}}, pre \rangle$ with $E_{+Gp} = \{e_p, e_{\bar{p}}, e_{\top}\}$, $R_a = (\{e_p, e_{\bar{p}}\} \times \{e_p, e_{\bar{p}}\}) \cup \{\langle e_{\top}, e_{\top} \rangle\}$ for $a \in G$ and $R_a = \{e_p, e_{\bar{p}}, e_{\top}\} \times \{e_{\top}\}$ otherwise, and the pre- and postconditions are defined as follows (see Figure 11):

- $pre_{e_p} = p$, $post_{e_p}(p) = 1$
- $pre_{e_{\bar{p}}} = \neg p$, $post_{e_{\bar{p}}}(p) = 0$
- $pre_{e_{\top}} = \top$

As before, we also use E_{+Gp}^* to denote the points of reference $\{e_p, e_{\bar{p}}\}$.

Private raising awareness could be used to raise awareness of a single agent when $G = \{a\}$, i.e. $+_{\{a\}}p$. In that case, only agent a has her awareness of p raised, and the other agents remain in the old state.

Because the pre- and postconditions ‘coincide’ (that is, the valuation of the postcondition verifies the precondition), we can draw the event model as in Figure 11, where written p inside an event e means that $pre(e) = p$ and $post_e(p) = 1$, and \top denotes $pre_e = \top$ and $post_e(p)$ is undefined for every $p \in P$.




Fig. 11 The event model \mathcal{E}_{+Gp} for private raising awareness of p , $+Gp$.

After privately raising the awareness of p , a formula ϕ holds if ϕ holds at the mappings $\langle w, e_p \rangle$ and $\langle w, e_{\bar{p}} \rangle$ in the product update, because the multi-pointed event model is defined as $\langle \mathcal{E}_{+Gp}, \{e_p, e_{\bar{p}}\} \rangle$. We also refer to this multi-pointed event model by $+Gp$ and the product update $\mathcal{M} \otimes \mathcal{E}_{+Gp}$ as \mathcal{M}^{+Gp} . We do not require this for e_{\top} because this event only represents that the other agents do not observe the event.

Like with public and private announcements, privately raising the awareness of p for the group of all agents, $G = \mathcal{A}$, amounts to raising (public) awareness of p .

Proposition 5 *Given a pointed ParEAL model $\langle \mathcal{M}, w \rangle$, a proposition p and a formula ϕ . Then $\mathcal{M}, w \models [+_{\mathcal{A}p}] \phi$ if and only if $\mathcal{M}, w \models [+p] \phi$ and $\mathcal{M}, w \models [+_{\mathcal{A}p}] \phi$ if and only if $\mathcal{M}, w \models [+p] \phi$.*

Proof The model $\mathcal{M}^{+\mathcal{A}p}$ consists of two disconnected components: the first one is exactly \mathcal{M}^{+p} , the result of taking the product update of \mathcal{M} with the part of $\mathcal{E}_{+\mathcal{A}p}$ which is identical to \mathcal{E}_{+p} (i.e. the events e_p and $e_{\bar{p}}$ and the relations between them), and the second component is the result of taking the product update of \mathcal{M} with the single event e_{\top} . The two components are disconnected because the edges between them, holding for $\mathcal{A} \setminus G$, are empty when $G = \mathcal{A}$.

[only if] Assume that $\mathcal{M}, w \models [+p]\phi$. Then $\forall \langle w, e \rangle \in W^{+p}$ s.t. $e \in E_{+p}^*$: $\mathcal{M}^{+p}, \langle w, e \rangle \models \phi$. But these worlds are exactly these worlds in the image of w under $+\mathcal{A}p$ that are used to determine satisfiability, and these worlds satisfy the same formulas. Therefore $\mathcal{M}, w \models [+_{\mathcal{A}p}]\phi$.

Similarly, assume that $\mathcal{M}, w \models [+p]\phi$. Then $\exists \langle w, e \rangle \in W^{+p}$ s.t. $e \in E_{+p}^*$: $\mathcal{M}^{+p}, \langle w, e \rangle \models \phi$. But again, this world is also in the image of w under $+\mathcal{A}p$ that is used to determine satisfiability, and this world satisfies the same formulas. Hence $\mathcal{M}, w \models [+_{\mathcal{A}p}]\phi$.

[if] We can reverse the reasoning above to show that $\mathcal{M}, w \models [+_{\mathcal{A}p}]\phi$ implies $\mathcal{M}, w \models [+p]\phi$ and $\mathcal{M}, w \models [+_{\mathcal{A}p}]\phi$ implies $\mathcal{M}, w \models [+p]\phi$ because satisfiability for $+\mathcal{A}p$ is determined by the worlds $\langle w, e_p \rangle$ and $\langle w, e_{\bar{p}} \rangle$ in the image of w under $+\mathcal{A}p$, because $e_{\top} \notin E_{+\mathcal{A}p}^*$. And these worlds are also in the image of w under $+p$ and again, satisfy the same conditions. So whenever something holds for all these worlds (for \models) or for at least one of the worlds (\models) under $+\mathcal{A}p$, this must also be the case for these worlds under $+p$.

Raising private awareness for the group of all agents, i.e. $G = \mathcal{A} = \{a_1, \dots, a_n\}$, is not equivalent to privately raising awareness for agent a_1 , then a_2 , then a_3 , ..., until a_n . I.e., $+_{\{a_1\}}p; \dots; +_{\{a_n\}}p$ is not the same as $+_{G}p$. This is because in the first case, neither agent observes that the other agents raise their awareness. Hence, even though each agent becomes aware of p , this is not common information—but this is the case for $+_{G}p$.

In Figure 12, an example is shown of applying the event model for raising private awareness for a single agent.




Fig. 12 The event model of $+_{\{b\}}p$, $\mathcal{E}_{+\{b\}p}$, applied to an epistemic model \mathcal{M} , on the left. In the product update (right), the aware cell $\|(w, e_p)\|_a$ is indicated by a dotted box and the aware cell $\|(w, e_p)\|_b$ by a dashed box.

7 Raising awareness without disclosing truth

By switching to partial valuations, raising awareness of a proposition p acts on models by extending the valuation function in such a way that p and $\neg p$ become equally plausible for the agents. This implies that in our semantics, unlike previous work on awareness where raising awareness also unveiled truth values [9, 14, 15, 18, 22], raising awareness does not disclose truth. This was illustrated in Figure 12 in which agent b becomes aware of p without knowing the truth value of p , and agent a does not change her mind. In this section, we provide a proof.

First, we consider the knowledge and beliefs of agents under raising awareness. We prove that no new knowledge or belief is acquired:

- everything agents previously knew or believed, is still known and believed after raising awareness (*knowledge and belief preservation*), and
- everything* known or believed after raising awareness, was already known or believed before (*knowledge and belief correspondence*).

In the second clause, there is a catch. This holds only for formulas not including the proposition awareness is raised of. Hence for p -free ϕ when the modality is $+_{GP}$. This is for the obvious reason that tautologies involving p , for example $p \vee \neg p$, or formulas of the form $p \rightarrow \phi$ where ϕ was previously known or believed, will also be known or believed after raising the awareness of p . The term ‘everything*’ is used to exclude these cases.

In fact, preservation and correspondence do not only apply to epistemic formulas but also to non-epistemic formulas: truth is also preserved and corresponds to the case of formulas not including the proposition awareness is raised of. This is straightforward: raising awareness extends the valuation function with proposition p but does not alter the valuation of other propositions. Therefore, formulas not involving p keep their old truth value, and formulas involving p acquire a truth value.

Proposition 6 (Truth preservation and correspondence) *Let \mathcal{M} be a ParEAL model and let w be a world in \mathcal{M} . Then for any group $G \subseteq \mathcal{A}$ and for any formula ϕ :*

$$\mathcal{M}, w \models \phi \rightarrow [+_{GP}]\phi \quad (\text{truth preservation})$$

and for any ψ not containing p :

$$\mathcal{M}, w \models [+_{GP}]\psi \rightarrow \psi \quad (\text{truth correspondence})$$

Proof For preservation, assume that $\mathcal{M}, w \models \phi$. Then, after raising the awareness of p , for all $\langle w, e \rangle \in W^{+_{GP}}$ with $e \in E_{+_{GP}}^*$ the valuation $V_{\langle w, e \rangle}^{+_{GP}}$ is either equal to V_w (in case $p \in \text{Dom}(V_w)$) or extends it V_w by a valuation for proposition p (in case $p \notin \text{Dom}(V_w)$). In both cases, for all $q \in \text{Dom}(V_w)$: $V_{\langle w, e \rangle}^{+_{GP}}(q) = V_w(q)$. Therefore, for all $\langle w, e \rangle \in W^{+_{GP}}$ with $e \in E_{+_{GP}}^*$ it holds that $\mathcal{M}^{+_{GP}}, \langle w, e \rangle \models \phi$, and thus $\mathcal{M}, w \models [+_{GP}]\phi$.

For correspondence, assume that $\mathcal{M}, w \models [+_{GP}]\psi$ for some formula ψ not containing p . This means that $\mathcal{M}^{+_{GP}}, \langle w, e \rangle \models \psi$ for all $\langle w, e \rangle \in W^{+_{GP}}$ with $e \in E_{+_{GP}}^*$. But since ψ did not contain p , for all q in ψ : $V_{\langle w, e \rangle}^{+_{GP}}(q) = V_w(q)$. Therefore, it must also hold that $\mathcal{M}, w \models \psi$.

In the proof of Proposition 6, there is no distinction between propositional formulas ϕ and non-propositional formulas ϕ , for example $\phi = K_a\psi$, because in both cases all the propositions q occurring in ϕ are evaluated at v . For the non-propositional formulas, this is true because of the specification property for valuations in ParEAL: this causes any propositions q occurring in ϕ such that $K_a\phi$, even if $\phi = K_b\psi$, to have a truth value at v . Then, the fact that raising awareness does not alter the valuation of these propositions q completes the proof.

Furthermore, raising the awareness of p does not disclose the truth of p itself. That is, whenever p was undefined ($\neq p$ and $\neq p$), it is false that after awareness is raised of p , p is true, and that p is false, i.e. $\models [+_{Gp}]p$ and $\models [+_{Gp}]\neg p$.

Proposition 7 (Raising awareness without disclosing truth) *Let \mathcal{M} be a ParEAL model and let w be a world in \mathcal{M} . Then for any group $G \subseteq \mathcal{A}$ and any proposition p :*

If $\mathcal{M}, w \neq p$ and $\mathcal{M}, w \neq p$ then $\mathcal{M}, w \models [+_{Gp}]p$ and $\mathcal{M}, w \models [+_{Gp}]\neg p$

Proof Whenever $\mathcal{M}, w \neq p$ and $\mathcal{M}, w \neq p$, it means that $p \notin \text{Dom}(V_w)$. Hence, $V_{\langle w, e_p \rangle}^{+_{Gp}}(p) = 1$ and $V_{\langle w, e_{\bar{p}} \rangle}^{+_{Gp}}(p) = 0$. Thus $\mathcal{M}^{+_{Gp}}, \langle w, e_p \rangle \models p$ and $\mathcal{M}^{+_{Gp}}, \langle w, e_{\bar{p}} \rangle \models \neg p$. By satisfiability for ParDEL+, then $\mathcal{M}^{+_{Gp}}, \langle w, e_p \rangle \models \neg p$. Thus, since $E_{+_{Gp}}^* = \{e_p, e_{\bar{p}}\}$, $\mathcal{M}, w \models [+_{Gp}]p$ and $\mathcal{M}, w \models [+_{Gp}]\neg p$.

In conclusion, raising awareness does not disclose truth values nor add new truths, other than the trivial cases ($p \vee \neg p$ or $p \rightarrow \phi$, where ϕ was true before, amongst others). This means that the raising awareness modalities introduced for ParEAL truly disconnect awareness from truth.

8 Forgetting awareness

We have defined a novel way to capture awareness of agents directly in the model structures using partial valuation functions and weakly reflexive relations. Moreover, we have introduced a dynamic modality for raising awareness on this semantics. A natural question to ask is then: can we define a reverse modality? I.e. is there a modality $\neg p$ for *forgetting awareness*?

Reverse modalities and operations have been studied throughout the history of logic: the AGM model for belief revision considers expansion as well as contraction [1], temporal logics are defined in function of *future* and *past* modalities [36,37] and in other work on awareness both raising and forgetting modalities have been introduced [9,14,15,18]. Besides the theoretical motivation, there is also a practical motivation to study forgetting coming from multi-agent systems. In such systems, agents use different ontologies and alignments to represent their knowledge and beliefs. During communication, they may encounter a counter-example to their alignments that they revise accordingly [39,40]. However, they do not need specific examples to communicate successfully, so they do not store them. In other words, they forget the examples. Therefore, a logical modeling of such agents could benefit from forgetting [10,11].

At first sight, forgetting may be considered as the opposite of raising awareness, reversing the models back to the initial situation before awareness was raised. Such

an operation can easily be achieved by reducing the domain of the valuation function by the proposition that is forgotten of and merging worlds up to bisimilarity: the proposition becomes undefined and hence, agents become unaware. This also forces models on which the awareness of p is raised and then directly forgotten to be bisimilar to the original situation, see Figure 13.




Fig. 13 Raising the awareness of p , $+p$, (left to right) and forgetting awareness of p , $-p$, (right to left), merging worlds up to bisimilarity. It holds that that $\mathcal{M}^{+p;-p}$ and \mathcal{M} are bisimilar.


We can also capture such a forgetting operation with event models for ParEAL, by extending the definition of a postcondition in Definition 27 with an option that we call \perp , as follows:

- $post : E \rightarrow (P \rightarrow \{0, 1, \perp\})$ is a *postcondition function* assigning to each event a partial function $post_e : P \rightarrow \{0, 1, \perp\}$.

Then, whenever for an event $e \in E$ and a proposition $p \in P$ it holds that $post_e(p) = \perp$, we “delete” the valuation of p in the product update. This means that in Definition 28, we adjust the clause for the new valuation function as follows:

$$- V_{\langle w, e \rangle}^{\mathcal{M} \otimes \mathcal{E}}(p) = \begin{cases} post_e(p) & \text{if } post_e(p) = 1 \text{ or } post_e(p) = 0 \\ \text{undefined} & \text{if } post_e(p) = \perp \\ V_w(p) & \text{otherwise} \end{cases}$$

The event model for forgetting the awareness of a proposition p would then consist of a single event with precondition \top and postcondition $post(p) = \perp$:



However, unlike raising awareness, this type of forgetting cannot be performed independently of forgetting truth. Recall that when raising their awareness, agents come to consider equally possible that a proposition is true and false—hence, they do not acquire any knowledge or belief about it. In contrast, when forgetting their awareness, agents also forget truth values (because they are no longer aware of a

proposition, they do not know whether it is true or not)—and hence, whatever was known or believed about it, is lost. This means that forgetting is not a true inverse operation of raising awareness, but more the reverse of raising awareness followed by an announcement of the proposition or its negation.

Yet, we may still study this forgetting modality, $-p$. For example, to investigate what happens if the proposition whose awareness is forgotten was used as evidence for another proposition. Consider the consecutive upgrade $+p; +q; !(p \rightarrow q); !p; -p$, where awareness of p and then of q is raised, a relation between the two propositions ($p \rightarrow q$) is established, p is announced and finally p is forgotten again. When $-p$ is applied, the truth value of q remains untouched: q will still be true even though its justification, p , is removed. But when we forget p , should we not also forget the conclusion q ?

There are two options. Either we can keep the conclusion q and view forgetting as a generalization or abstraction modality getting rid of the evidence but keeping the conclusions, or we can discard the conclusions as well to arrive back at the original state (unless, of course, the conclusions are supported independently from p —in the example before, if we have both $p \rightarrow q$, $r \rightarrow q$ and r , we can still deduce q when forgetting p). For both options, there are arguments: for full logical agents, forgetting evidence should imply to forget its conclusions since the deductions that led to them can no longer be repeated, whereas for non-logical agents, it may not be necessary to keep all evidence that led them to certain conclusions because these conclusions alone may be sufficient for them to communicate successfully, like in [20, 21]. Therefore this form of forgetting could benefit a formal model of non-logical agents [10, 11].

A definitive answer about the different views on forgetting is beyond the scope of this paper, but awareness does provide a starting point for the discussion.

9 Conclusion

We have introduced a new semantics for agent awareness using partial valuation functions and weakly reflexive relations. This semantics allows agents to use different signatures to describe their knowledge and beliefs that other agents may be unaware of. Whenever agents communicate, what they learn is described by two steps: (1) raising their awareness of the propositions they hear, and (2) learning the truth values of these propositions in the form of knowledge or belief through announcements, radical and conservative upgrades. Unlike other approaches to extend DEL with agent awareness, our work completely disconnects awareness from truth: awareness is raised without affecting the knowledge and beliefs of agents.

As a consequence of this two-step approach to learning ((1) raising awareness and (2) learning truth), models are able to truly dynamically evolve: any future evolution of agents' knowledge, belief and awareness is not bound by the initial state. In contrast, in DEL, any evolution was bound and agents can only learn by restricting their possibilities: eliminating worlds or reorganizing them in a different way with respect to the plausibility relations. These features are also available in our framework, as the dynamic upgrades from DEL remain, but now worlds may see their scope increasing through raising awareness. This open approach to modeling the evolution of agents' epistemic states is desirable for dynamic and

open multi-agent systems where agents are required to continuously adapt to their environment or to information learned from other agents [2, 20, 21].

The effect of dropping reflexivity to capture agent awareness and unawareness also opens the door for interpreting complex epistemic upgrades, in particular those that involve a component of privacy. Private announcements have been introduced for DEL in [3], however, they do not preserve the typical structure of DEL models: reflexivity is lost in order to ensure privacy. As a solution, it has been suggested that the only dynamic upgrades permitted are those that preserve the properties of the DEL models [3]. This means that private announcements are no longer permitted to be performed, let alone more complex dynamic upgrades involving privacy. The work we present in this paper offers a different solution: by weakening the properties of the models, i.e. replacing reflexivity by weak reflexivity, we do not only provide a novel way to approach agent awareness, but also allow for more complex upgrades involving privacy to be interpreted. To illustrate this, private announcements do preserve weak reflexivity. As such, our framework can be seen as a more general framework for modelling communication and knowledge or belief change in multi-agent systems.

Dropping reflexivity also has a drawback: factivity no longer holds. However, it still holds locally within the aware cells of agents. This corresponds to the weak reflexivity condition required for the accessibility relations. As a result, knowledge may be considered ‘intermediate’ between $S5$ (within aware cells) and $KD45$ (outside aware cells). It is an open question to what axiom such an intermediate knowledge operator corresponds that formalizes weak reflexivity as a condition on the models.

In DEL, reduction axioms are used to ‘reduce’ formulas with dynamics, such as announcements, to a formula without dynamics recursively. This is used to reduce expressivity, soundness and completeness to the case of modal logic. In our framework, these reduction axioms do not translate directly. This is because preconditions are used in a different way: they are used to select worlds that “do not falsify the precondition”, i.e. they either verify the precondition, or the precondition is undefined. Since undefinedness cannot be captured in our language, this complicates reducing formulas with raising awareness modalities to formulas without them. A three-valued approach may be useful for this purpose and could be used to define an axiomatization of our logic.

This work can be extended to investigate a dual-operator for raising awareness: forgetting. The discussion whether forgetting should serve as an abstraction operator, which experiments have shown does not negatively affect communication between adaptive agents [20, 21], or should function more truthfully to logical agents could be a starting point for this purpose. Furthermore, another notion of forgetting may be considered that preserves awareness while deleting the truth of propositions. Such a forgetting modality would cause agents to unlearn what they knew or believed, while their awareness is unaffected.

ParDEL+ is an extension of DEL in the sense that if all agents were aware of the same propositions and would not use raising awareness, the logics would coincide. In addition, with ParDEL+ one can consider situations starting from no awareness and raising agent awareness progressively to reach a DEL state. Hence, DEL with public signature awareness may be considered as the limit of ParDEL+ restricted to raising only the DEL signature. Yet, it may not be strictly necessary

to reach full signature awareness for agents to communicate satisfactorily with each other: they only need to be aware of the subsignatures of other agents that occur in their communication. It would be worth investigating this relation further.

Acknowledgements This work has been partially supported by MIAI @ Grenoble Alpes (ANR-19-P3IA-0003).

Conflict of interest

The authors declare that they have no conflict of interest.

References

1. Alchourrón, C.E., Gärdenfors, P., Makinson, D.: On the logic of theory change: Partial meet contraction and revision functions. *The Journal of Symbolic Logic* **50**(2), 510–530 (1985)
2. Atencia, M., Schorlemmer, W.M.: An interaction-based approach to semantic alignment. *J. Web Semant.* **12**, 131–147 (2012). DOI 10.1016/j.websem.2011.12.001
3. Baltag, A., Moss, L.S., Solecki, S.: The logic of public announcements and common knowledge and private suspicions. In: *Proceedings of the 7th Conference on Theoretical Aspects of Rationality and Knowledge*, Evanston, USA, 1998, pp. 43–56 (1998)
4. Baltag, A., Smets, S.: Dynamic belief revision over multi-agent plausibility models. In: *Proceedings of the 2016 conference on Logic and the Foundations of Game and Decision Theory*, vol. 6, pp. 11–24. University of Liverpool (2006)
5. Baltag, A., Smets, S.: Group belief dynamics under iterated revision: fixed points and cycles of joint upgrades. In: *Proceedings of the 12th Conference on Theoretical Aspects of Rationality and Knowledge*, pp. 41–50. ACM (2009)
6. van Benthem, J.: Dynamic logic for belief revision. *Journal of Applied Non-Classical Logics* **17**(2), 129–155 (2007). DOI 10.3166/jancl.17.129-155
7. van Benthem, J.: *Logical dynamics of information and interaction*. Cambridge University Press (2011)
8. van Benthem, J., van Eijck, J., Kooi, B.: Logics of communication and change. *Information and Computation* **204**(11), 1620–1662 (2006)
9. van Benthem, J., Velázquez-Quesada, F.R.: The dynamics of awareness. *Synthese* **177**(1), 5–27 (2010)
10. van den Berg, L.: *Cultural knowledge evolution in dynamic epistemic logic*. Ph.D. thesis, University of Grenoble (2021)
11. van den Berg, L., Atencia, M., Euzenat, J.: A logical model for the ontology alignment repair game. *Autonomous Agents and Multi-Agent Systems* **35**(2), 1–34 (2021)
12. Blackburn, P., De Rijke, M., Venema, Y.: *Modal Logic*. Cambridge Tracts in Theoretical Computer Science. Cambridge University Press (2001). DOI 10.1017/CBO9781107050884
13. Brouwer, L.E.J.: *Over de grondslagen der wiskunde*. Maas & van Suchtelen (1907)
14. van Ditmarsch, H., French, T.: Awareness and forgetting of facts and agents. In: *2009 IEEE/WIC/ACM International Joint Conference on Web Intelligence and Intelligent Agent Technology*, vol. 3, pp. 478–483. IEEE (2009)
15. van Ditmarsch, H., French, T.: Becoming aware of propositional variables. In: *Indian Conference on Logic and Its Applications*, pp. 204–218. Springer (2011)
16. van Ditmarsch, H., French, T., Velázquez-Quesada, F.R.: Action models for knowledge and awareness. In: *Proceedings of the 2012 International Conference on Autonomous Agents and Multi-Agent Systems*, pp. 1091–1098 (2012)
17. van Ditmarsch, H., French, T., Velázquez-Quesada, F.R., Wáng, Y.N.: Knowledge, awareness, and bisimulation. *arXiv preprint arXiv:1310.6410* (2013)
18. van Ditmarsch, H., Herzig, A., Lang, J., Marquis, P.: Introspective forgetting. *Synthese* **169**(2), 405–423 (2009)
19. van Ditmarsch, H., van der Hoek, W., Kooi, B.: *Dynamic Epistemic Logic*, vol. 337. Springer Science & Business Media (2007)

20. Euzenat, J.: First experiments in cultural alignment repair (extended version). In: *The Semantic Web: ESWC 2014 Satellite Events - ESWC 2014 Satellite Events*, Anissaras, Crete, Greece, May 25-29, 2014, Revised Selected Papers, pp. 115–130 (2014). DOI 10.1007/978-3-319-11955-7_10
21. Euzenat, J.: Interaction-based ontology alignment repair with expansion and relaxation. In: *Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence, IJCAI 2017, Melbourne, Australia, August 19-25, 2017*, pp. 185–191 (2017). DOI 10.24963/ijcai.2017/27
22. Fagin, R., Halpern, J.Y.: Belief, awareness, and limited reasoning. *Artificial Intelligence* **34**(1), 39–76 (1987)
23. Fagin, R., Moses, Y., Halpern, J.Y., Vardi, M.Y.: *Reasoning about knowledge*. MIT press (2003)
24. Gärdenfors, P., Rott, H., Gabbay, D., Hogger, C., Robinson, J.: Belief revision. *Computational Complexity* **63**(6), 35–132 (1995)
25. Halpern, J.Y.: Alternative semantics for unawareness. *Games and Economic Behavior* **37**(2), 321–339 (2001)
26. Halpern, J.Y., Rêgo, L.C.: Reasoning about knowledge of unawareness. *Games and Economic Behavior* **67**(2), 503–525 (2009)
27. Hansen, J.U.: Modeling truly dynamic epistemic scenarios in a partial version of del. *The Logica Yearbook 2013* pp. 63–75 (2014)
28. Heifetz, A., Meier, M., Schipper, B.C.: Interactive unawareness. *Journal of economic theory* **130**(1), 78–94 (2006)
29. Hill, B.: Awareness dynamics. *Journal of Philosophical Logic* **39**(2), 113–137 (2010)
30. Hintikka, J.: Knowledge and belief: An introduction to the logic of the two notions. *Studia Logica* **16** (1962)
31. van der Hoek, W., Jaspars, J., Thijsse, E.: Honesty in partial logic. *Studia Logica* **56**(3), 323–360 (1996)
32. Jaspars, J., Thijsse, E.: *Fundamentals of partial modal logic*. Studies in Logic Language and Information (1996)
33. Kleene, S.C.: *Introduction to metamathematics* (1952)
34. Modica, S., Rustichini, A.: Awareness and partitioned informational structures. In: *Epistemic Logic and the Theory of Games and Decisions*, pp. 151–168. Springer (1997)
35. Plaza, J.: Logic of public communications. In: Z.W. Ras (ed.) *Proceedings of the 4th International Symposium on Methodologies for Intelligent Systems*, pp. 201–216. North-Holland (1989)
36. Prior, A.: Time and modality. *Revue Philosophique de la France Et de l* **148** (1957)
37. Prior, A.: Past, present and future. *Revue Philosophique de la France Et de l* **157** (1967)
38. Rott, H.: Conditionals and theory change: Revisions, expansions, and additions. *Synthese* **81**(1), 91–113 (1989)
39. Shvaiko, P., Euzenat, J.: Ontology matching: state of the art and future challenges. *IEEE Transactions on Knowledge and Data Engineering* **25**(1), 158–176 (2013)
40. Tamma, V., Cranefield, S., Finin, T.W., Willmott, S.: *Ontologies for agents: Theory and experiences*. Springer Science & Business Media (2005)
41. Thijsse, E.: *Partial logic and knowledge representation* (1994)
42. Veltman, F.: Defaults in update semantics. *Journal of philosophical logic* **25**(3), 221–261 (1996)